



This is a repository copy of *On the application of generative adversarial networks for nonlinear modal analysis*.

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/182989/>

Version: Accepted Version

Article:

Tsialiamanis, G., Champneys, M.D. orcid.org/0000-0002-3037-7584, Dervilis, N. orcid.org/0000-0002-5712-7323 et al. (2 more authors) (2022) On the application of generative adversarial networks for nonlinear modal analysis. *Mechanical Systems and Signal Processing*, 166. 108473. ISSN 0888-3270

<https://doi.org/10.1016/j.ymssp.2021.108473>

© 2021 Elsevier Ltd. This is an author produced version of a paper subsequently published in *Mechanical Systems and Signal Processing*. Uploaded in accordance with the publisher's self-archiving policy. Article available under the terms of the CC-BY-NC-ND licence (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

On the application of generative adversarial networks for nonlinear modal analysis

G. Tsialiamanis^{1a}, M.D. Champneys^a, N. Dervilis^a, D.J. Wagg^a, K. Worden^a

^a*Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield
Mappin Street, Sheffield S1 3JD*

Abstract

Linear modal analysis is a useful and effective tool for the design and analysis of structures. However, a comprehensive basis for *nonlinear* modal analysis remains to be developed. In the current work, a machine learning scheme is proposed with a view to performing nonlinear modal analysis. The scheme is focussed on defining a one-to-one mapping from a latent ‘modal’ space to the natural coordinate space, whilst also imposing orthogonality of the mode shapes. The mapping is achieved via the use of the recently-developed *cycle-consistent generative adversarial network* (cycle-GAN) and an assembly of neural networks targeted on maintaining the desired orthogonality. The method is tested on simulated data from structures with cubic nonlinearities and different numbers of degrees of freedom, and also on data from an experimental three-degree-of-freedom set-up with a column-bumper nonlinearity. The results reveal the method’s efficiency in separating the ‘modes’. The method also provides a nonlinear superposition function, which in most cases has very good accuracy.

Keywords: Generative adversarial networks (GANs), cycleGAN, nonlinear modal analysis, inductive biases.

1. Introduction

Many approaches have been followed throughout the years in order to perform dynamic analysis of structures, the most dominant being *modal analysis* [1]. The reason that modal analysis has proved so powerful is because it provides a meaningful *decomposition* of the oscillations of a structure. The components of this decomposition, the *modes*, are independent, and each one has its own dynamic characteristics, i.e. mode shape, natural frequency, modal damping, modal mass, etc; providing the users with independent

¹Corresponding Author: George Tsialiamanis (g.tsialiamanis@sheffield.ac.uk)

movements to study, and ways to isolate potential problems while designing
10 a structure. For linear structures, linear modal analysis is by far the most
common method of analysis [2]. Both in modelling and at the experimental
level, modal decomposition of linear structural behaviour has been achieved
to a high level of accuracy. The scheme also provides a practical solution for
15 multi-degree-of-freedom (MDOF) systems by decoupling them into single-
degree-of-freedom (SDOF) systems, using an eigen-decomposition of the
system matrices where the modes are related to eigenvectors and natural
frequencies related to the eigenvalues.

Apart from assisting in understanding the resonant behaviour of structures,
modal analysis also provides a convenient model reduction technique. Higher-
20 frequency modes absorb less energy from the excitation and so tend to affect
the behaviour of the structure less; thus, they can quite often be omitted.
Under a such a framework, techniques, such as principal component analysis
(PCA) [3] and proper orthogonal decomposition (POD) [4], are similarly
followed with a view to identifying vibration modes from available data [5] and
25 to solve by projecting the algebraic systems onto a lower-dimensional vector
basis. The aforementioned scheme is followed when one has no information
about the excitation and is often referred to as *operational modal analysis*
(OMA) or *data-driven* modal analysis.

Despite the fact that existing modal analysis methods are focussed on linear
30 systems, they are often used on structures with suspected nonlinearities.
Modal analysis of structures, which may have structural elements with
nonlinear behaviour, is often plausible in cases of real structures, for some
range of external loads that does not suffice for the nonlinearities to affect
the behaviour of the structure. Consider the *Duffing oscillator* [6], which
35 has a cubic term in its differential equation. For small values of excitation
force, the system may not exhibit notable nonlinear behaviour, making
linear analysis appropriate and sufficient. For MDOF systems with similar
nonlinearities, decomposition of the movement into modes may be achieved
in similar cases of low-force excitation, if there is a stable underlying linear
40 system; a necessary but not always sufficient criterion. Nevertheless, this is
not always the case. Real-life structures exhibit nonlinear behaviour quite
often and linear modal analysis methods fail to define nonlinear modes of
vibration. Methods have been developed to deal with such issues [7], but only
achieve preservation of a *subset* of the properties of linear modal analysis
45 [8, 9].

An approach to data-driven nonlinear modal analysis using machine learning
was proposed in [10] and [11]. The idea of structural independence, together
with a Shaw-Pierre concept [12] ansatz were used to motivate a new definition
of nonlinear normal modes (NNMs). In [10], a genetic algorithm was used
50 to define a decomposition of the displacements of various nonlinear systems

into a modal space. With the objective function of the algorithm, a type of orthogonality of the modes, via their statistical independence, was enforced. The approach agreed with the POD in the limit of linear behaviour. The results were encouraging and revealed that such an approach can decouple
55 NNMs from nonlinear systems more effectively than conventional linear decomposition methods. Although the approach presented in [10] yielded a good modal decomposition for a two-degree-of-freedom simulated system with cubic nonlinearity, the results were not as good for a three-degree-of-freedom system with similar nonlinearity and for an experimental case. Moreover,
60 the best model selection was achieved ‘by eye’, since the objective function of the optimisation problem did not yield the best results in terms of mode separation. Also, the nonlinear superposition function studied in [10] did not yield satisfactory results.

In the current work, following the same machine learning framework and
65 defining an NNM using the same assumptions, an alternative approach is followed. Instead of trying to decompose the natural coordinate space into modal coordinates, a mapping between the two is sought, using a predefined modal space. In order to achieve the desired result, a recently-developed algorithm is used, the *cycle-consistent generative adversarial network* (cycle-
70 GAN) [13]. The algorithm is used to define a forward mapping from the natural coordinate space to a modal space, as well as the inverse mapping, to achieve (nonlinear) superposition of the modes. A great advantage of the algorithm, as it will be described, is its invertibility property, so one has, as a consequence, a smooth mapping from one space to the other. Furthermore, a
75 second neural network assembly is used that forces the transition from modal to natural coordinates to enforce the orthogonality of the mode shapes and so to satisfy an NNM criterion. Compared to previous approaches, the new method is a way of training both forward and backward mappings at the same time and identifying the mapping that most efficiently separates the
80 structural movement into independent modes.

The layout of the paper is as follows. Section 2 provides a brief introduction to generative adversarial networks, to problems that arise when using them and to the cycle-GAN algorithm and how it resolves some of these problems. In Section 2.4, the layout of the neural network, which is used in parallel with
85 the cycle-GAN in order to maintain the desired orthogonality, is described. In Section 3 the proposed nonlinear modal analysis algorithm is described. In Section 4, applications to simulated dynamical systems are presented, together with an application to an experimental dataset. Section 5 considers the inverse mapping and superposition of the modes. In Section 6 the
90 correlation of the modes is studied for two correlation metrics, a linear and a nonlinear metric. Finally, in Section 7, conclusions are drawn about the method.

2. Generative adversarial networks (GAN)

2.1. Vanilla GAN

95 A recent approach to machine learning is that of using *generative models*. Such models are able to learn the distribution of given data and generate artificial data according to it. A trivial approach to the problem of generating artificial data would be to define a Gaussian distribution using the mean and the covariance of the data. Following such an approach, the normal
100 distribution could be used to generate artificial data. However, real data distributions are more complicated; for example, to cope with multi-modal distributions, a kernel density estimate [14] could be calibrated according to the data.

A more recent approach is to use a *Generative Adversarial Network* (GAN)
105 [15]. The algorithm was initially created to generate synthetic images that look real; i.e. the model learns how to embed figures into some latent space and simultaneously how to generate data according to a proper distribution. Apart from the main goal of the algorithm, a novel way of training neural networks was introduced. *Adversarial training* is defined as a competition
110 between two neural networks. In the basic GAN, the first network is the *generator*, which tries to generate samples that look real and the second is the *discriminator*, which tries to identify whether a sample comes from the real dataset or is artificial.

Training is orchestrated as a competition between two networks. The *dis-*
115 *criminator* D is a network with an output representing the probability of its input sample \mathbf{x} being real; i.e. $P_{\mathbf{x} \sim p_{data}} = D(\mathbf{x})$. Throughout training, real and fake samples are introduced to the discriminator and using back-propagation it becomes better at distinguishing samples from the real dataset from generated/artificial samples. The discriminator essentially draws a
120 decision boundary around the manifold of the available data. On the other hand, the *generator* G takes as input a noise vector \mathbf{z} from some pre-defined probability distribution $p_z(\mathbf{z})$ and creates a sample $G(\mathbf{z})$ in the feature space of the dataset. Thereafter, the sample is passed through the discriminator in order to decide whether it is real or generated. The probability of a
125 generated sample being real is given by $D(G(\mathbf{z}))$. Forcing the generator to create samples that ‘fool’ the discriminator into classifying them as real (i.e. minimisation of $\log(1 - D(G(\mathbf{z})))$), results in creating samples/images that look real. The optimisation problem based on an objective function $V(D, G)$ for the training of both networks is given by,

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))] \quad (1)$$

130 where $\mathbb{E}_{x \sim p_{data}}[\cdot]$ is the mean value of the functions within the brackets $[\cdot]$ with respect to the distribution p_{data} .

The layout of a basic GAN is shown in Figure 1. In practice, training is performed in two steps. During the first step, a batch of samples is randomly sampled from the dataset along with an equally-large batch sampled from the generator. Training of the discriminator *only* is then performed using as target labels 1 (real) for the dataset samples and 0 (fake) for the generated samples. During this step, the discriminator is trained to better distinguish, real from fake samples. During the second step, the generator is trained, while the parameters of the discriminator are held constant. For this step, a batch of noise vectors are sampled and the two networks are connected together as shown in Figure 1. The second term of equation (1) is used alone for training and the target labels for the output of the discriminator are 1s, meaning that the generator should transform the noise vectors into samples that the discriminator accepts as real.

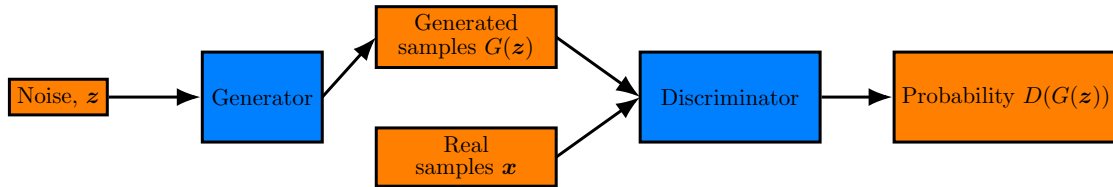


Figure 1: Layout of a basic (vanilla) GAN in full assembly; the generator transforms noise into generated samples and the discriminator attempts to distinguish between real and generated samples.

145 GANs in structural dynamics can be used in order to generate artificial data; in cases where acquiring data from structures is expensive, this aspect of GANs is useful and reduces the cost of recording data. GANs may also be used in other ways, to learn mappings from one space to another. A simple approach would be to learn a mapping from a latent space to a high-dimensional space of data. One would hope that via such mappings, each latent variable would encode distinct features of the data, but this is not usually the case; GANs tend to provide an entangled representation of the data via their latent variables.

In order to illustrate the problem of the aforementioned entangled representation, a linear example borrowed from *structural health monitoring* (SHM) is presented. A GAN is used to map a two-dimensional Cartesian space to the manifold formed by collecting *frequency response function* (FRF) samples of a simulated system under various levels of damage. Given a lumped-mass system as an example (Figure 2), and introducing two damage cases, one for spring 1 and one for spring 2, samples of the FRF of the first degree of freedom can be collected (damage is simulated as a stiffness reduction [16]). Performing a PCA on the dataset in order to visualise the data, the first three

principal component scores of the dataset are shown in Figures 3a and 3b as the blue semiopaque points. The variations applied were stiffness reductions for the first and second spring in the interval $[0\%, 20\%]$ and combinations according to the Cartesian product $[0\%, 20\%] \times [0\%, 20\%]$.

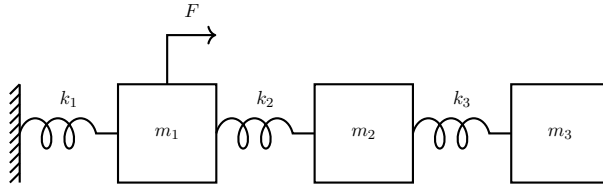
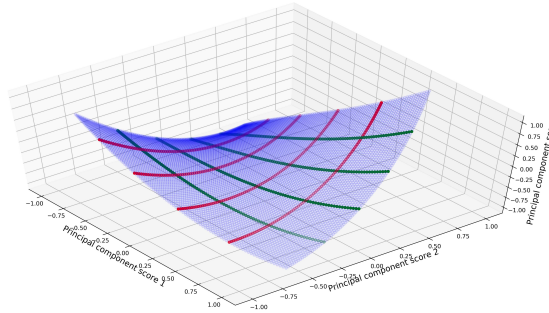
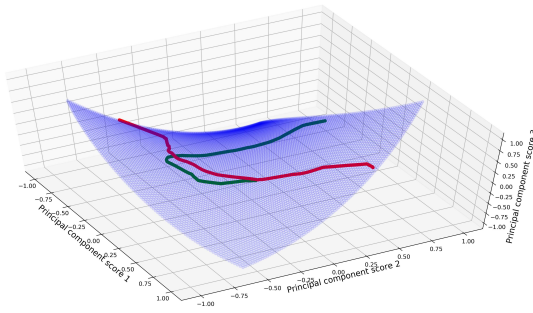


Figure 2: Three-degree-of-freedom mass spring system.



(a)



(b)

Figure 3: First three principal component scores of the dataset (blue). On the top, points corresponding to constant damage for spring 1 and varying damage percentage for spring 2 (red), and points corresponding to constant damage for spring 2 and varying damage percentage for spring 1 (green). In the bottom, GAN-generated samples (red and green) by locking one latent variable to 0 and varying the other in the interval $[-1, 1]$; regular GAN used.

In Figure 3a, curves are shown, along which one of the damage parameters is constant and the second varies. There are two damage parameters and no other parameter varies during simulation, so the manifold is two-dimensional and each curve is parametrised by one of the two damage parameters.

After collecting the data, they are used to train a GAN. The noise vector is a two-dimensional vector and the latent variables are drawn from uniform distributions over the interval $[-1, 1]$ and the output space is the three-dimensional principal component space of the collected FRFs. As a universal GAN cross-validation scheme does not exist, the GAN was trained according to the procedure followed in the original paper [15]. Using the trained generator to output data by varying one latent variable at a time and keeping the other constant, the resulting samples are shown in Figure 3b. It is clear that a disentangled representation of the manifold has not been achieved, since the green and red lines intersect. The algorithm is not expected to perform so, since there are no restrictions on how the latent variables are used in order to generate artificial data. To achieve a disentangled representation in the current work, the inductive bias of *invertibility* is enforced in the mapping provided by the GAN.

2.2. Invertibility of the GAN

In an attempt to avoid entanglement, as shown in Figure 3b, imposition of the property of invertibility to a GAN is considered in the current work. Given that a mapping is invertible, meaning that it is a bijective mapping from one manifold to another, entanglement in Figure 3b would probably be alleviated and the mappings should be smoother and more meaningful.

Given a continuous bijection $\phi : M \rightarrow N$, where M and N are manifolds, it is straightforward to show that if c_1, c_2 are curves in N , then $\phi^{-1} \circ c_1$ and $\phi^{-1} \circ c_2$ are curves in M with the same number of points of intersection. Considering $\phi : M \rightarrow N$ and two curves $c_1, c_2 \in N$ then $\phi^{-1} \circ c_1 = c'_1$ and $\phi^{-1} \circ c_2 = c'_2$ are curves in M . For every parameter t_i corresponding to the intersections of the two curves c_1, c_2 , it stands that $c_1(t_i) = c_2(t_i)$, which leads to $c'_1(t_i) = c'_2(t_i)$ because ϕ is a bijection. Therefore, c_1 and c_2 have the same points of intersection as c'_1 and c'_2 . Consequently, by introducing the invertibility inductive bias in a GAN, helps to avoid mappings like the one shown in Figure 3b, where two orthonormal axes from the latent variable space are mapped onto the entangled green and red lines. Enforcing invertibility might force the GAN to more efficiently achieve a disentanglement of the features of the data in the, not so rare, case that underlying parameters affect the behaviour of the system in a different way (as in the SHM example presented above) and the mapping from the parameter space to the feature space is a bijection.

In order to visualise the effects of imposing the property of invertibility on

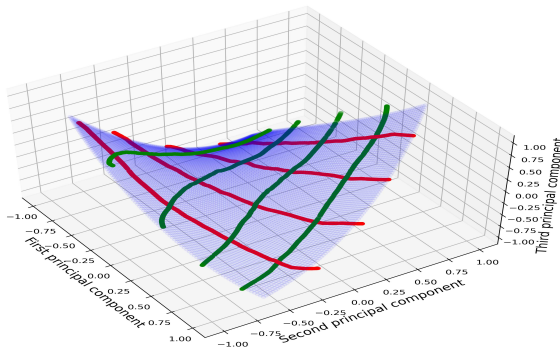


Figure 4: First three principal component scores of the dataset (blue) and cycleGAN-generated samples (red and green) by locking one latent variable to various constant values and varying the other the interval $[-1, 1]$.

training of the GAN, the cycleGAN algorithm, which will be described in detail, is applied on the same dataset. Similar plots of the results are shown
 210 in Figure 4. It is clear that there is no entanglement of features present, in contrast to when training a regular GAN without any restrictions and that the mapping from the two-dimensional noise to the three-dimensional feature space is closer to the mapping shown in Figure 3a.

For modal analysis, invertibility is quite important. Firstly, it ensures a
 215 bijective mapping, and therefore a unique encoding of every different state of the structure into a modal space. Secondly, securing an inverse mapping is fundamental, since a way to perform superposition of the modes is desirable. To satisfy the aforementioned criteria, the algorithm chosen herein is the *cycle-consistent generative adversarial network* [13]. In the following sub-
 220 sections, this variation of the GAN is presented, together with a way to induce the desired orthogonality of the modes in the mapping.

2.3. Cycle-GAN

The algorithm chosen to search for invertible mappings under the framework of GANs, is that of the *Cycle-Consistent Adversarial Network* (cycle-GAN)
 225 [13]. The algorithm was initially introduced as an attempt to transfer images from one style to another; for example, to transform photos into paintings in the style of famous artists. What is interesting, is the way this is achieved; the whole procedure is similar to defining an autoencoder [17], but in terms of a GAN.

230 A cycle-GAN uses two generators ($G_{X \rightarrow Y}$ and $G_{Y \rightarrow X}$) and two discriminators (D_X and D_Y). Each pair of networks acts similarly to the classic GAN scheme. The layout of a cycle-GAN is shown in Figures 5 and 6. The first generator

learns to transform samples from domain X into samples of domain Y . If used on images, the samples most probably live in the same space, but in different regions or sub-manifolds. The second generator learns to map figures from domain Y back to domain X . More specifically, as shown in Figure 5, the generator $G_{Y \rightarrow X}$ learns to map the sample generated by $G_{X \rightarrow Y}$ back to the original sample in domain X . Training is performed in two steps; during the first step, the procedure from domain X to domain Y is followed (Figure 5) and during the second, the opposite (Figure 6).

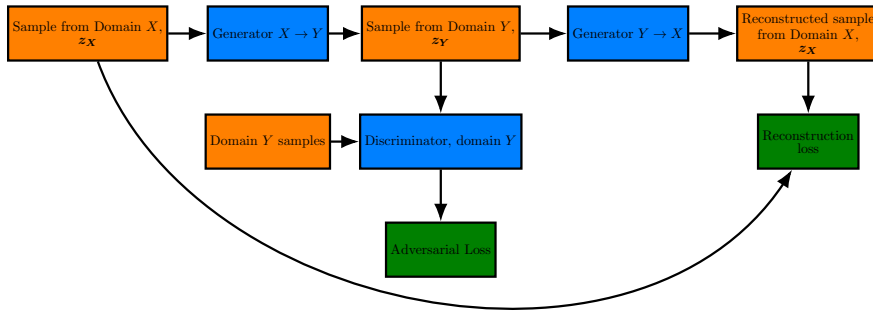


Figure 5: Cycle GAN layout assembled in order to learn the mapping from domain X to domain Y and back. Samples are converted from X to Y by the generator $G_{X \rightarrow Y}$. The generated samples are used for adversarial training of the discriminator D_Y and the generator $G_{X \rightarrow Y}$. Subsequently, the samples are inverse-mapped back to domain X via generator $G_{Y \rightarrow X}$ and both generators are trained using the reconstruction-loss error.

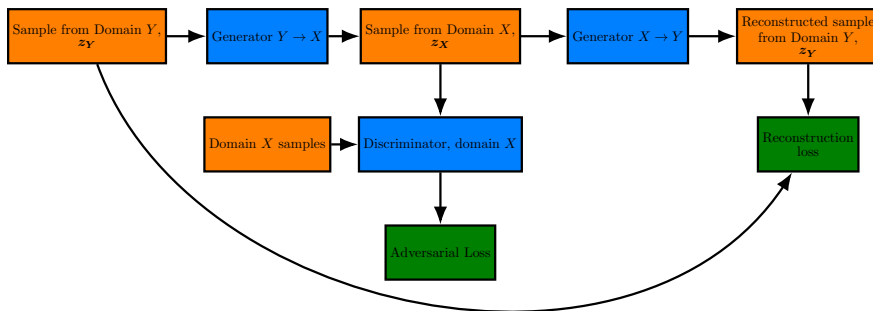


Figure 6: Cycle GAN layout assembled in order to learn the mapping from domain Y to domain X and backwards. Anti-symmetrical to the procedure shown in Figure 5.

In both steps, the *adversarial loss* \mathcal{L}_1 is computed, exactly as in the GAN

scheme from,

$$\begin{aligned} \mathcal{L}_1(G_{X \rightarrow Y}, D_Y, X, Y) = & \mathbb{E}_{\mathbf{y} \sim p_y(\mathbf{y})} [\log D_Y(\mathbf{y})] + \\ & \mathbb{E}_{\mathbf{x} \sim p_x(\mathbf{x})} [\log(1 - D_Y(G_{X \rightarrow Y}(\mathbf{x})))] \end{aligned} \quad (2)$$

and accordingly for the inverse training step. The second type of loss used in training - *the reconstruction loss (cycle loss in the original work)* - is given by,

$$\begin{aligned} \mathcal{L}_2(G_{X \rightarrow Y}, G_{Y \rightarrow X}) = & \mathbb{E}_{\mathbf{x} \sim p_x(\mathbf{x})} [\|G_{Y \rightarrow X}(G_{X \rightarrow Y}(\mathbf{x})) - \mathbf{x}\|_n] + \\ & \mathbb{E}_{\mathbf{y} \sim p_y(\mathbf{y})} [\|G_{X \rightarrow Y}(G_{Y \rightarrow X}(\mathbf{y})) - \mathbf{y}\|_n] \end{aligned} \quad (3)$$

where $\|\cdot\|_n$ is the n^{th} order norm. In the original work, a first-order norm was used, but in the current work, a second-order one yielded better results.

The total training loss is computed from,

$$\begin{aligned} \mathcal{L}(G_{X \rightarrow Y}, G_{Y \rightarrow X}, D_X, D_Y) = & \mathcal{L}_1(G_{X \rightarrow Y}, D_Y, X, Y) + \\ & \mathcal{L}_1(G_{Y \rightarrow X}, D_X, Y, X) + \\ & \lambda \mathcal{L}_2(G_{X \rightarrow Y}, G_{Y \rightarrow X}) \end{aligned} \quad (4)$$

where λ controls the relative importance between the adversarial loss and the reconstruction loss (in the original work, a nominal value suggested is $\lambda = 10$). The optimisation problem solved is,

$$G_{X \rightarrow Y}^*, G_{Y \rightarrow X}^* = \min_{G_{X \rightarrow Y}, G_{Y \rightarrow X}} \max_{D_X, D_Y} \mathcal{L}(G_{X \rightarrow Y}, G_{Y \rightarrow X}, D_X, D_Y) \quad (5)$$

In practice, training is performed in the two stages described by Figures 5 and 6, and each stage comprises three training steps. The first two steps are similar to the GAN scheme described in Section 2.1. The newly-introduced
 245 third step is that of the reconstruction loss. Samples that change domain in each stage via each generator, are mapped back to their original domain and the reconstruction loss is computed. The error is back-propagated and both generators' trainable parameters are calibrated according to it.

250 Using this scheme instead of a *vanilla* GAN ensures the invertibility of the mappings and therefore the advantages discussed in the previous section. The discriminators, for the purposes of the current work, have an auxiliary role in the training procedure.

2.4. Orthogonality enforcement

255 Orthogonality constraints have been used in GANs before [18], with a view to dealing with mode (not in a structural modal analysis sense) collapse problems [19]. Orthogonality may be considered as another inductive bias that users impart in training, in order for the results to be closer to their

physical understanding of the data. In some cases, orthogonality may assist
 260 in achieving disentanglement of features. In the current work, orthogonality
 is a desired property of the modal analysis procedure. As described in
 the next sections, the approach to be followed, using cycle GANs, assumes
 domain X to be the natural coordinate space of the displacements of some
 structure and domain Y the modal space. Under this framework, it is desired
 265 that samples mapped from modal to natural coordinates be orthogonal when
 they correspond to different modes, enforcing the orthogonality of the mode
 shapes.

Departing for a while from the modal framework and returning to GANs,
 orthogonality enforcement means that “two vectors tangent to the latent
 270 manifold and parallel to two of the axes of the latent space are, of course,
 orthogonal and shall remain orthogonal in the real space. A schematic
 representation of the idea is shown in Figure 7. Trying to enforce this
 behaviour in the generator, a new assembly of networks is defined with a
 view to locally ensuring the orthogonality of the partial derivatives of the
 275 mappings from the modal space to the real space. Mappings with such
 behaviour are called *conformal* or *angle-preserving*.

The layout of the assembly used to enforce orthogonality is schematically
 shown in Figure 8. The goal is to maintain orthogonality of the grid of the
 latent space into the real space. As shown in Figure 8, a random latent
 280 point \mathbf{u}_1 is sampled. An axis/variable in the latent space is then chosen,
 and a small quantity ϵ is added to that latent coordinate to get the point
 $\mathbf{u}_{1a}^+ = \mathbf{u}_1 + \{0, 0, \dots, \epsilon, 0, \dots, 0\}$. In the same way, but by subtracting ϵ , yields
 the point $\mathbf{u}_{1a}^- = \mathbf{u}_1 - \{0, 0, \dots, \epsilon, 0, \dots, 0\}$. Afterwards, another axis/latent
 variable is chosen and the same procedure is repeated generating points
 285 $\mathbf{u}_{1b}^+ = \mathbf{u}_1 + \{0, 0, \dots, 0, \epsilon, 0, \dots, 0\}$ and $\mathbf{u}_{1b}^- = \mathbf{u}_1 - \{0, 0, \dots, 0, \epsilon, 0, \dots, 0\}$.

All these points are passed through the generator and their real space
 counterparts are generated ($\mathbf{y}_{1a}^+ = G(\mathbf{u}_{1a}^+)$, $\mathbf{y}_{1a}^- = G(\mathbf{u}_{1a}^-)$, $\mathbf{y}_{1b}^+ = G(\mathbf{u}_{1b}^+)$
 and $\mathbf{y}_{1b}^- = G(\mathbf{u}_{1b}^-)$). Consequently, the vectors \mathbf{v}_{1a} and \mathbf{v}_{1b} are computed as
 $\mathbf{v}_{1a} = \mathbf{y}_{1a}^+ - \mathbf{y}_{1a}^-$ and $\mathbf{v}_{1b} = \mathbf{y}_{1b}^+ - \mathbf{y}_{1b}^-$. Finally, the inner product between
 290 \mathbf{v}_{1a} and \mathbf{v}_{1b} is computed and divided by the quantity $\|\mathbf{v}_{1a}\| \|\mathbf{v}_{1b}\|$ to get
 the cosine of the angle between the two vectors. Optimising the Generator’s
 weights so that this quantity is as close to zero as possible will enforce
 the desired orthogonality. The orthogonality assembly proposed is in fact
 calculating the numerical gradient (in geometrical terms $\frac{\partial}{\partial u_i}$). along two
 295 different axes of a specific point in the manifold in the real space. In the
 case of modal analysis, as it will be described, if this bias is enforced in the
 generator mapping from modal to natural coordinates, the orthogonality of
 the mode shapes is ensured.

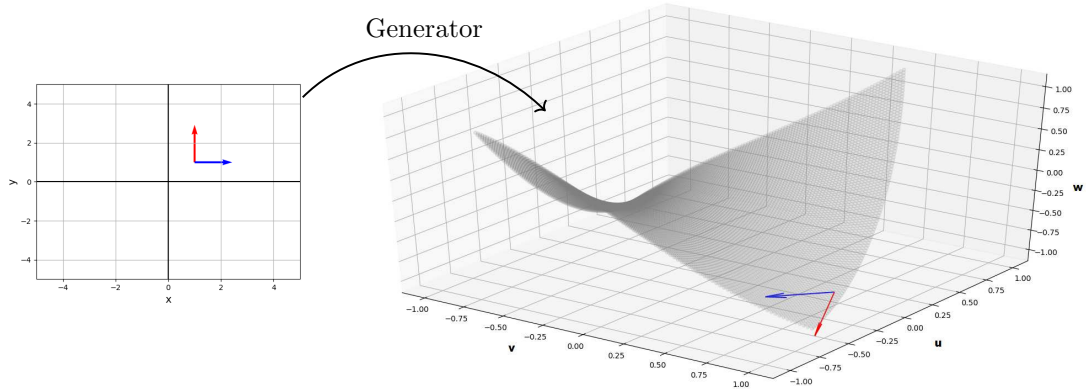


Figure 7: Preservation of orthogonality of vectors by the mapping of the generator from the source Cartesian space (x, y coordinate system) to the target manifold (u, v, w coordinate system).

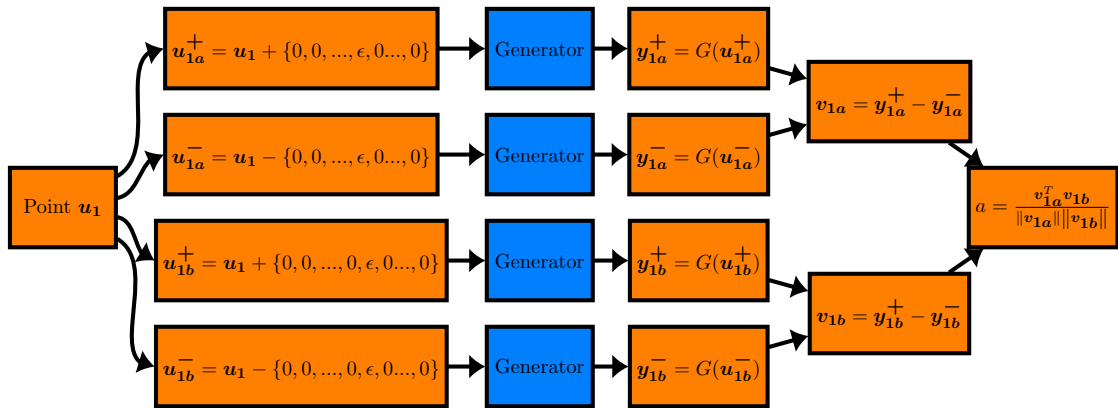


Figure 8: Orthogonality enforcement assembly using the generator that maps samples from the source Cartesian space to the target manifold.

3. Performing nonlinear modal analysis using a cycle-GAN

300 As described, the cycle-GAN provides a practical method for creating bijective mappings from a source domain to a target domain; this is useful for modal analysis. Given a linear structure, mode shapes provide a representative way of analysing the various independent ways (*modes*) that the structure oscillates under some external load. The modes in such cases, can be used
 305 to decompose any response of the structure into independent responses, each one referring to a different natural frequency. Unfortunately this is not the case for nonlinear structures.

In linear structures, modal analysis can be performed either by eigenvalue analysis of the structural parameters (mass and stiffness matrices), or in
 310 an operational manner, via *principal component analysis* PCA [3] of the displacements (or accelerations) given by sensors placed on a structure; this procedure in some cases coincides with linear modal analysis. For structures with nonlinearities, these methods cannot be applied. PCA itself is a linear method and therefore provides only a linear decomposition of the data.
 315 Motivated by PCA's data-driven scheme, an attempt to perform similarly data-driven nonlinear modal analysis was proposed in [10] and [11].

The decomposition in [10] aimed at maintaining the statistical independence and orthogonality aspects of a modal analysis. More specifically, a genetic algorithm was used to learn a decomposition of the displacements into latent
 320 variables with correlation close to 0. The approach performed better than a linear PCA analysis of the data and, under the criteria of [10, 11], in the case of a two-degree-of-freedom lumped-mass system with a cubic nonlinearity, a very efficient decomposition into modes was achieved. However, as the degrees-of-freedom of the systems increase, the algorithm seems to not
 325 perform equally well, because the mapping was more complicated.

A major drawback of the algorithm is that the maps used in order to perform the decomposition are restricted to a fixed polynomial order. For example, for a three-degree-of-freedom system, the equation used to transform the physical coordinates $\{y_1, y_2, y_3\}^T$ into the corresponding modal $\{u_1, u_2, u_3\}^T$ was,

$$\begin{Bmatrix} u_1 \\ u_2 \\ u_3 \end{Bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{Bmatrix} y_1 \\ y_2 \\ y_3 \end{Bmatrix} + \begin{bmatrix} b_{11} & \dots & a_{19} \\ b_{21} & \dots & b_{29} \\ b_{31} & \dots & a_{39} \end{bmatrix} \begin{Bmatrix} y_1^3 \\ y_1^2 y_2 \\ y_1^2 y_3 \\ y_2^3 \\ y_2^2 y_1 \\ y_2^2 y_3 \\ y_3^3 \\ y_3^2 y_1 \\ y_3^2 y_2 \end{Bmatrix} \quad (6)$$

where $a_{ij}, i, j \in \{1, 2, 3\}$ and $b_{i,j}, i \in \{1, 2, 3\}, j \in \{1, 2, \dots, 9\}$ are the tunable parameters that are optimised using the objective function,

$$\begin{aligned}
J = & |\{A_1\} \cdot \{A_2\}| + |\{A_1\} \cdot \{A_3\}| + |\{A_2\} \cdot \{A_3\}| \\
& + \text{Cor}(u_1, u_2) + \text{Cor}(u_1, u_3) + \text{Cor}(u_2, u_3) \\
& + \text{Cor}(u_1^3, u_2) + \text{Cor}(u_1^3, u_3) + \text{Cor}(u_2^3, u_1) \\
& + \text{Cor}(u_2^3, u_3) + \text{Cor}(u_3^3, u_1) + \text{Cor}(u_3^3, u_2)
\end{aligned} \tag{7}$$

where $A_i = \{a_{1i}, a_{2i}, a_{3i}\}$. The problem with the decomposition of equation (6) is that it is a truncated polynomial, which is not a universal approximator.

The above objective function, aims at the orthogonality of the modal coordinates and their statistical independence in the u space. In the current work a different approach is followed. Instead of decomposing the dataset of displacements into a modal space, a mapping from a pre-defined latent space to the observed dataset and back will be sought. Using the cycle-GAN scheme, and considering domain X as the displacement domain and domain Y as the modal domain, this composition is attempted. For the rest of the paper, in order to use the same symbols as in [10], the physical domain and its corresponding coordinates will be Y and \mathbf{y}_i respectively and the ‘modal space’ and coordinates will be U and \mathbf{u}_i respectively.

For the purposes of modal analysis and trying to impose the statistical independence of the modal coordinates \mathbf{u}_i , they are chosen to be n -dimensional random Gaussian vectors, with mean values equal to zero and correlation matrix equal to I_n which is an n -dimensional identity matrix. Using this modal space, all correlation terms described as in equation (7) are expected to be zero.

To help ensure the statistical independence of the modal coordinates, PCA was applied to the physical coordinates before training the cycle-GAN. PCA is selected in order to minimise the initial correlations and because it is a linear and easily-invertible transformation.

In addition to minimising linear correlations between modal coordinates, PCA is expected to assist further. Assuming that the system’s response can be decomposed in modes which are excited on a different level and contribute unevenly to the total movement, a scaling problem arises. The axes of the modal space correspond to curves in the physical space. These curves, in the linear case, are the (linear) axes of the ellipsoid, which is formed if one plots points that correspond to displacements of the structure’s degrees-of-freedom in the physical displacement space. By introducing nonlinearities to the system, the aforementioned ellipsoid is deformed and its axes are no longer lines, they become curves. The algorithm proposed is called to find a mapping ϕ , which maps the axes of the predefined modal coordinates

onto the deformed-by-the-nonlinearity axes of the aforementioned ellipsoid.
 360 Because every mode contributes on a different level to the total motion of
 the structure, some of these axes are longer than others. By performing
 PCA, the coordinate system is rotated so that the axes' variance are in
 descending order. Subsequently, by scaling every coordinate of the samples
 in the interval $[-1, 1]$, as needed for the neural networks to perform, it is
 365 expected that the axes representing each mode will be scaled almost similarly,
 facilitating the task of the cycle-GAN; i.e. defining a mapping from the modal
 coordinates to the physical coordinates.

By optimising with a target value of zero for the cosine computed in the
 rightmost block of computation in Figure 8, distinct modal coordinates
 370 are forced to generate samples in the physical space that are orthogonal
 to the ones generated by different mode variables; this way imposing the
 orthogonality of the mode shapes and maintaining locally the orthogonality
 of the aforementioned curved axes in the real space. The major advantage
 of the proposed cycle-GAN approach is that, since the neural networks
 375 concerned can approximate any function [20], it is not restricted by the order
 of the terms used in equation (6). Furthermore, the inverse mapping provides
 the nonlinear superposition mapping required for the sake of completeness of
 the method. In the next sections, applications of the algorithm on simulated
 and experimental data are presented.

380 4. Case studies

As in [10], a two-degree-of-freedom system was studied first. In every sim-
 ulated case study here, the physical system is a lumped-mass system as
 shown in Figure 9. The parameters of the model are the same as in [10], i.e.
 $m = 1.0$, $c = 0.1$, $k = 10$, $k_3 = 1500$. The equations of motion of the system
 are,

$$\begin{bmatrix} m & 0 \\ 0 & m \end{bmatrix} \begin{Bmatrix} \ddot{y}_1 \\ \ddot{y}_2 \end{Bmatrix} + \begin{bmatrix} 2c & -c \\ -c & 2c \end{bmatrix} \begin{Bmatrix} \dot{y}_1 \\ \dot{y}_2 \end{Bmatrix} + \begin{bmatrix} 2k & -k \\ -k & 2k \end{bmatrix} \begin{Bmatrix} y_1 \\ y_2 \end{Bmatrix} + \begin{Bmatrix} k_3 y_1^3 \\ 0 \end{Bmatrix} = \begin{Bmatrix} F \\ 0 \end{Bmatrix} \quad (8)$$

The excitation was Gaussian white noise with zero mean and standard
 deviation 5.0, low-pass filtered onto the frequency interval $[0, 50]$ Hz. For
 every system, two datasets of 100000 points were generated using differently-
 seeded random excitations. The first dataset was used for training and model
 385 selection and the second for testing the efficiency of the algorithm, as well as
 for the *power spectral density* (PSD) functions presented. The PSDs were
 calculated using Welch's method [21].

For every case study, the generators used for the cycle-GAN were three-
 layered neural networks, since they are proven universal approximators [22].
 390 All tested neural networks had an input and an output layer, both with

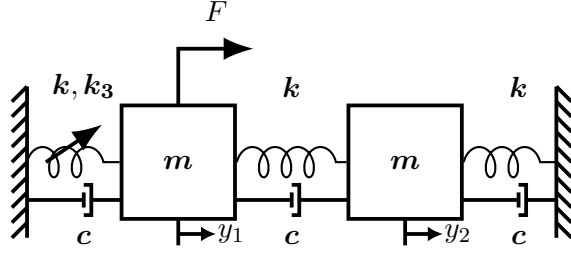


Figure 9: Two degree-of-freedom mass-spring system.

neurons equal to the dimension of the problem and a hidden layer whose size was optimised by training networks with different sizes and picking the best according to an inner product criterion described in the next paragraphs. The hidden layer sizes used belonged to the set $\{50, 60, \dots, 190, 200\}$. For every size, 20 random initialisations of the neural networks were performed. The large number of initialisations is determined by the obvious sensitivity of the algorithm to the initial placements of the principal axes on the target manifold. Moreover, for every generator hidden layer size, the corresponding discriminator had the same hidden layer size. This was done in order to allow both networks to act on the same latent space and in order to reduce the hyperparameters by inducing such a symmetry to the algorithm. The activation functions that performed best were *hyperbolic tangent* for the hidden layer and linear for the output layer, which is actually the best practice for regression problems. As proposed in [13], the value of the λ parameter from equation (4) is set to 10 and the training algorithm was the Adam optimiser [23].

In each case study, the best model was selected as the one that minimised the average cosine between every combination of PSD vectors in the latent space ((PSD_1, PSD_2) for the two-dimensional case, (PSD_1, PSD_2) , (PSD_1, PSD_3) , (PSD_2, PSD_3) for the three-dimensional, etc.). The equation describing the model selection criterion is defined by,

$$\mathcal{L}_{\cos} = \sum_{i=1, j=i+1}^{n_{dof}} \frac{PSD_i \cdot PSD_j}{\|PSD_i\| \|PSD_j\|} \quad (9)$$

where n_{dof} is the number of degrees of freedom of the system. The cosine criterion is an attempt to imitate the ‘by eye’ way of choosing the best model based on the isolation of modes in the PSDs. It is not identical, since minimisation of the dot product could be achieved if every PSD was equal to zero for every frequency, except for one; however, it is expected that this will not be the case, because of the orthogonality restriction applied. Model efficiency testing was performed every 100 epochs of training. The quantity in equation (9) cannot go to zero, since damping forces some of the energy

415 in the PSD to concentrate in frequencies around the natural frequencies;
nevertheless, it is a convenient measure of separation of the modes.

Throughout the rest of the paper, in order to keep a similar format to [10],
PSDs referring to physical coordinates will be drawn with a blue colour,
PSDs referring to modal components (or PCA components for the case of the
420 experimental data) will be drawn using a black colour and PSDs referring
to latent components of the cycle-GAN approach, presented herein, will be
drawn in a red colour.

4.1. Two-degree-of-freedom system

The described procedure was initially followed for a two-degree-of-freedom
425 simulated lumped mass system with a cubic nonlinearity (Figure 9). The
PSDs of the physical degrees of freedom are shown in Figure 10. The
nonlinearity is clearly affecting the PSD_1 . A spreading towards higher
frequencies is clear in the second mode. Linear natural frequencies are the
ones at which the structure absorbs the most energy and are proportional to
430 the stiffness of the structural members. As the nonlinear member in this case
is hardening, there are time instants during its movement that its stiffness
is higher because of the higher value of displacement of the first degree of
freedom, leading to the natural frequencies and energy in the PSD being
spread towards higher frequencies. The movement in this case is clear if one
435 considers the natural frequencies of the underlying linear problem, which are
0.5 Hz and 0.87 Hz.

Using the cycle-GAN to decompose the displacements into modal coordinates,
the model that yielded the best results had 100 units in its hidden layer.
The effect of the decomposition performed by that model is shown in the
440 bottom PSDs in Figure 11. For comparison, the PCA decomposition, which
in this case coincides with linear modal analysis, is shown on the top row
of the same figure. As expected, linear modal analysis cannot decouple the
modes, because of the nonlinearity. In contrast, the cycle-GAN algorithm is
able to efficiently do that. The result is very similar to that using the SADE
445 algorithm in [10].

The results of this case study, which can be considered a benchmark, are
encouraging. The large size of the hidden layer might be a result of selecting
the best model, having as the single criterion, the inner product in equation
(9). Smaller networks might also perform well. If one is concerned about the
450 complexity of the model, a second criterion can be considered along with
the inner product for the procedure of model selection. The algorithm is
subsequently tested on larger systems and also on an experimental system.

4.2. Three-degree-of-freedom system

In order to test the method on a three-degree-of-freedom system, a third
455 mass was added to the system shown in Figure 9, connected to the second

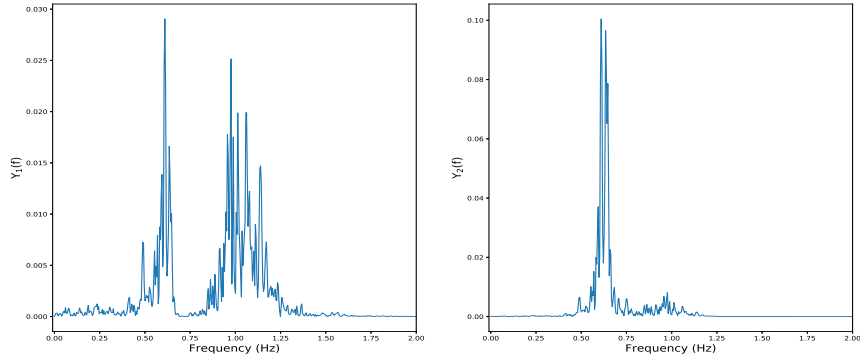


Figure 10: PSDs of two-degree-of-freedom structure; physical coordinates.

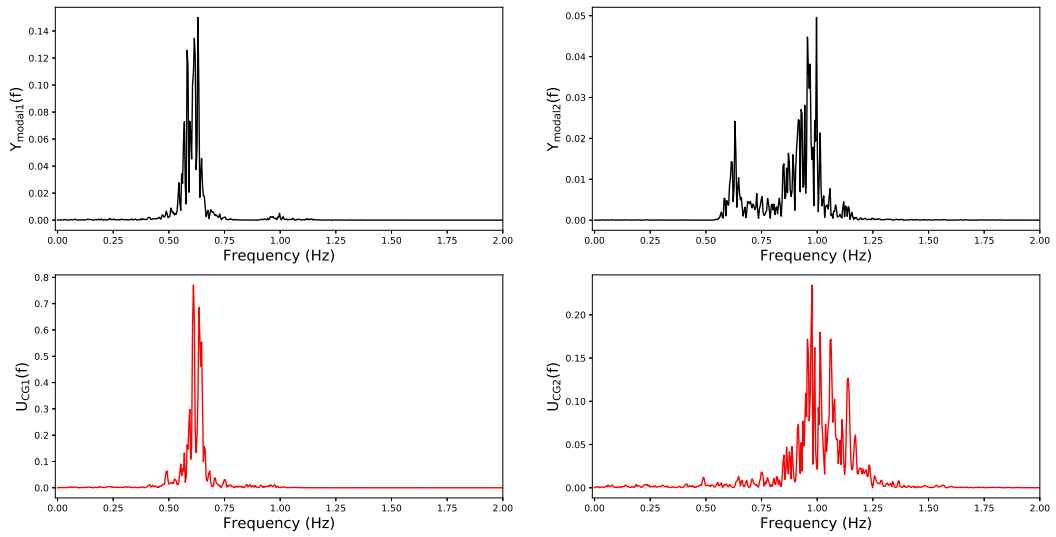


Figure 11: PSDs of two-degree-of-freedom structure, linear modal decomposition (top) and cycle-GAN decomposition selected via the inner product criterion (bottom).

mass and the ground with a linear spring with $k = 10$ and $c = 0.1$. Exactly the same procedure was followed. The PSDs for the physical coordinates are shown in Figure 12. The natural frequencies of the underlying linear system are 0.39 Hz, 0.71 Hz and 0.93 Hz and again the spread and the movement because of the nonlinearity are seen. The best model for the current case study had 110 hidden units and the decomposition provided by the algorithm in terms of PSDs is shown in the bottom row of Figure 13. Again, for comparison, the linear modal/PCA coordinate PSDs are shown in the top row. The decoupling of the modes is clear. In every PSD a different peak is dominant and only small effects of other modes are present.

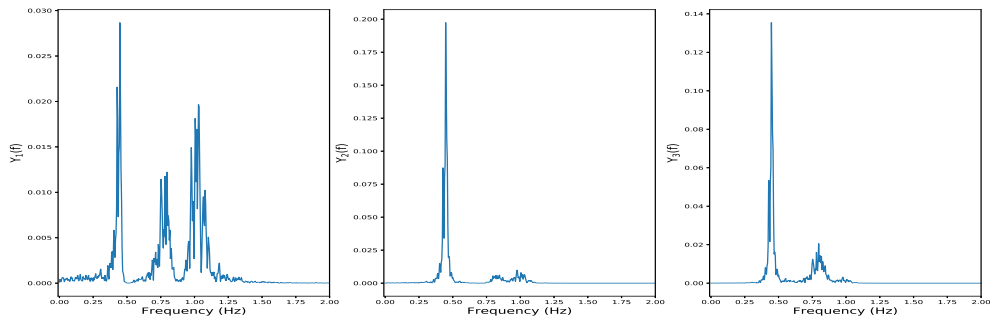


Figure 12: PSDs of three-degree-of-freedom structure; physical coordinates.

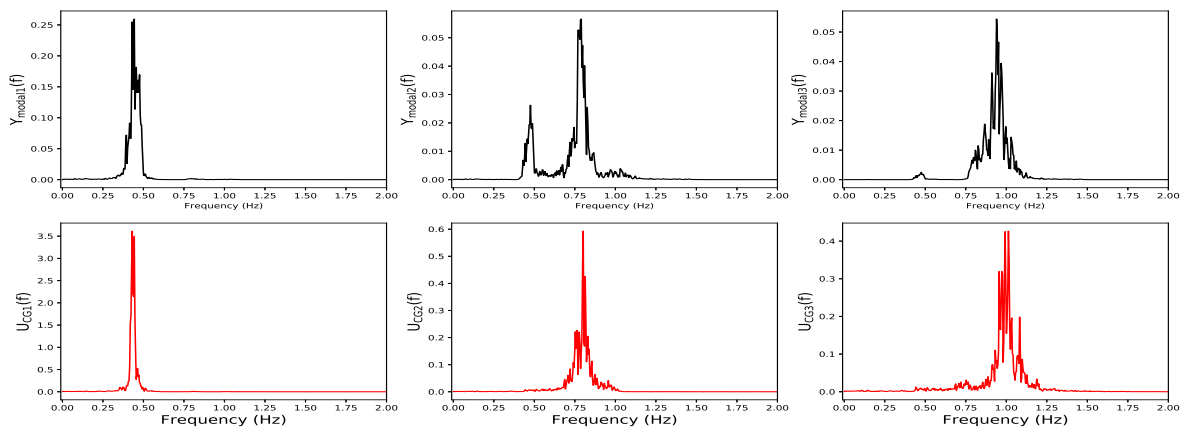


Figure 13: PSDs of three-degree-of-freedom structure, linear modal decomposition (top) and cycle-GAN decomposition selected via the inner product criterion (bottom).

4.3. Four-degree-of-freedom system

For the case of a four-degree-of-freedom system (which was not studied in [10]), it was noted that a larger value for the nonlinear stiffness parameter k_3 was needed in order for the structure to exhibit strongly-nonlinear behaviour; therefore, k_3 was increased to 3000 for this case study. In this case the model that had the best performance was a neural network with 100 neurons in its hidden layer. The PSDs of the four displacements are shown in Figure 14. The natural frequencies of the underlying linear system are 0.31 Hz, 0.59 Hz, 0.81 Hz and 0.96 Hz but again a movement and a spreading towards higher frequencies is observed. In the first PSD, four peaks are clearly seen. It is expected from the decomposition to have single peaks for every modal coordinate PSD. Using a linear modal decomposition, single peaks are not achieved, as seen in Figure 15 in the PSDs of the third and the fourth modal coordinates. However, using the proposed algorithm, in Figure 16, a single peak can be seen as dominant in every modal coordinate PSD.

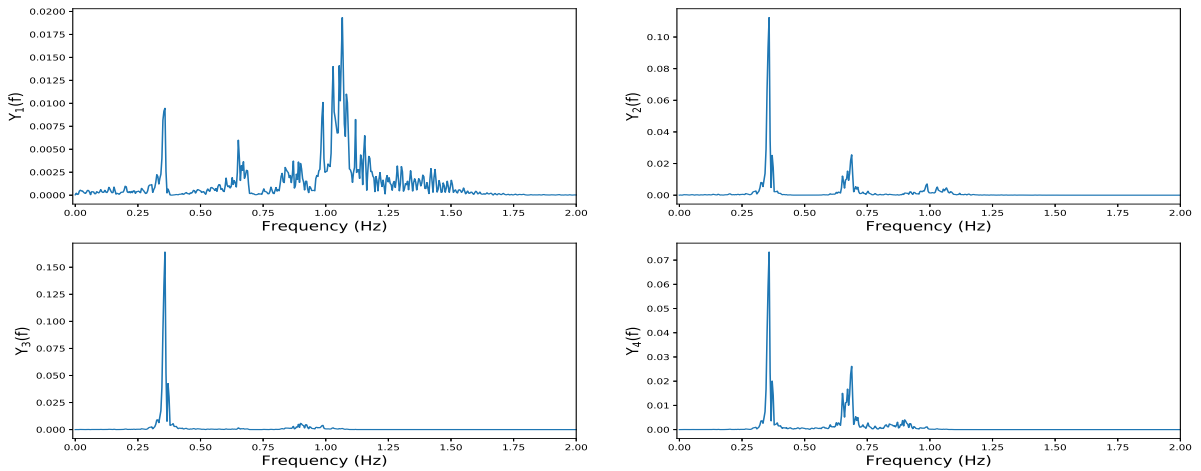


Figure 14: PSDs of four-degree-of-freedom structure; physical coordinates.

4.4. Three-degree-of-freedom experimental system

The experimental set-up for the data used in the current work was the one described in [24]; the experimental structure is shown in Figure 17. The structure was tested in 17 different states some of them being considered damaged and some not. The different states and their description are shown in Table 1. Particularly interesting are the states 10-14. In these states, a column between the second and the third floor is placed near a bumper, introducing a nonlinearity into the system in the form of bilinear stiffness. Each state has a different value for the gap between the column and the

490 bumper (from 0.20mm to 0.5mm). For the analysis here, the data from State 12 were used, where the gap was 0.13 mm, as well as from State 14, where the gap was 0.05 mm. The latter case was severely nonlinear because of the more frequent collisions. These cases may be considered harshly nonlinear compared to the smooth nonlinearities of the simulated case studies. The structure had four sensors recording accelerations, one on the base and one on each floor. For the current work, the displacements of the three floors are used, considering the base acceleration as the excitation (as in an earthquake).

495

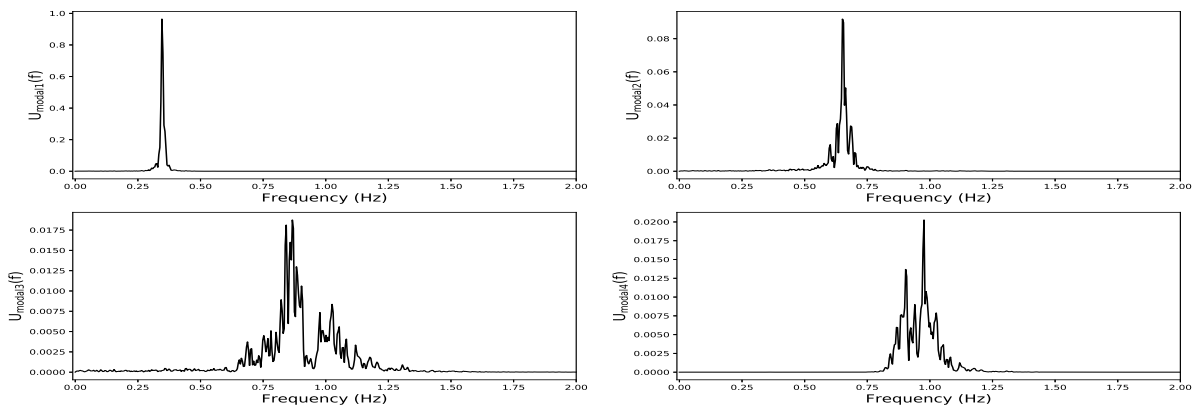


Figure 15: PSDs of four-degree-of-freedom structure, linear modal decomposition coordinates.

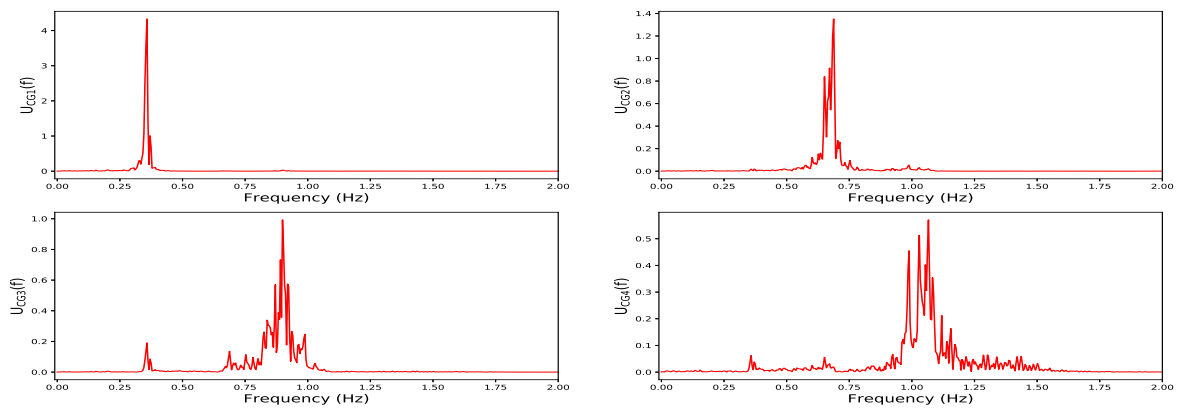


Figure 16: PSDs of four-degree-of-freedom structure selected via the inner product criterion, cycle-GAN latent variables.

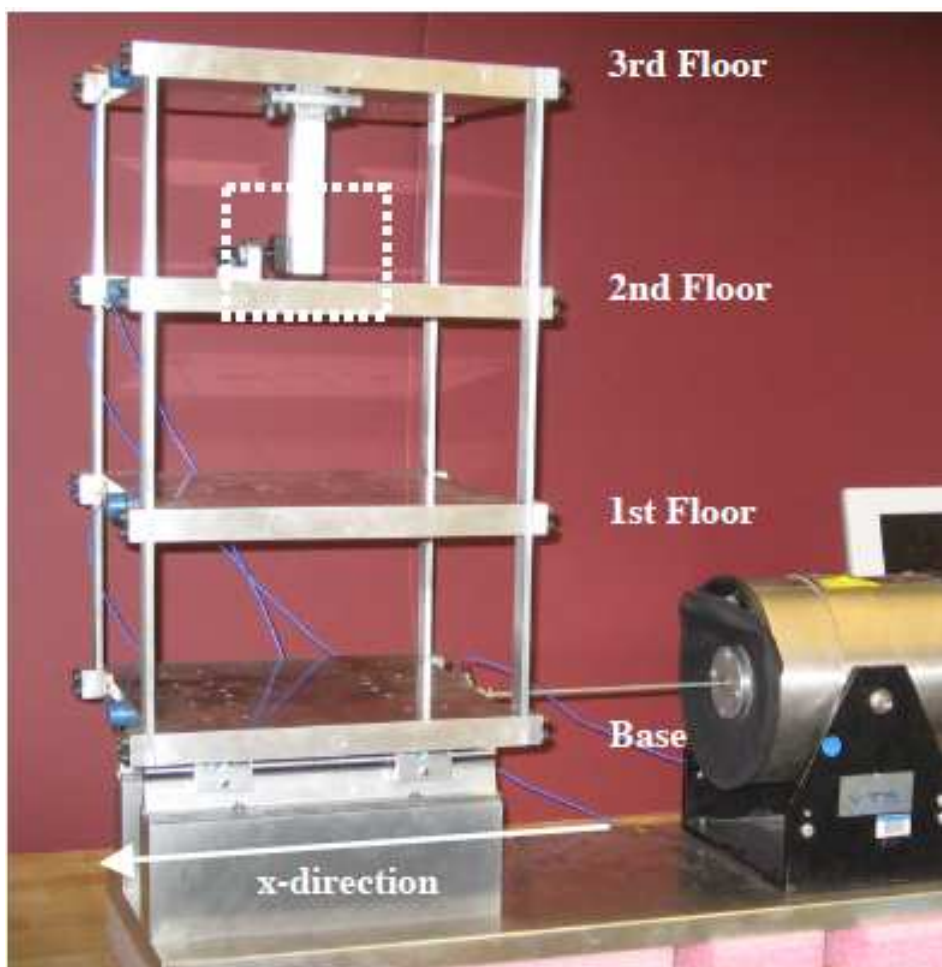


Figure 17: Experimental set-up of three-floor and the bumper nonlinearity between the second and third floor shown in the dashed box [24].

Label	State Condition	Description
State #1	Undamaged	Baseline condition
State #2	Undamaged	Added mass (1.2 kg) at the base
State #3	Undamaged	Added mass (1.2 kg) on the 1 st floor
State #4	Undamaged	Stiffness reduction in base front column
State #5	Undamaged	Stiffness reduction in base front and rear column
State #6	Undamaged	Stiffness reduction in 1 st floor front column
State #7	Undamaged	Stiffness reduction in 1 st floor front and rear column
State #8	Undamaged	Stiffness reduction in 2 nd floor front column
State #9	Undamaged	Stiffness reduction in 2 nd floor front and rear column
State #10	Damaged	Gap (0.20 mm)
State #11	Damaged	Gap (0.15 mm)
State #12	Damaged	Gap (0.13 mm)
State #13	Damaged	Gap (0.10 mm)
State #14	Damaged	Gap (0.05 mm)
State #15	Damaged	Gap (0.20 mm) and mass (1.2 kg) at the base
State #16	Damaged	Gap (0.20 mm) and mass (1.2 kg) at the 1 st floor
State #17	Damaged	Gap (0.10 mm) and mass (1.2 kg) at the 1 st floor

Table 1: Description of different states of the experimental set-up [24].

500 *4.4.1. Experimental data: State 12*

Similar to the simulated examples, the PSDs of the signals of the three sensors are shown in Figure 18. The PSDs in the specific figure are computed from one out of 50 experiments performed corresponding to State 12. Three peaks can clearly be seen. For the rest of the experimental case studies presented, the PSDs are calculated as the average PSD of the 50 experiments performed.

505

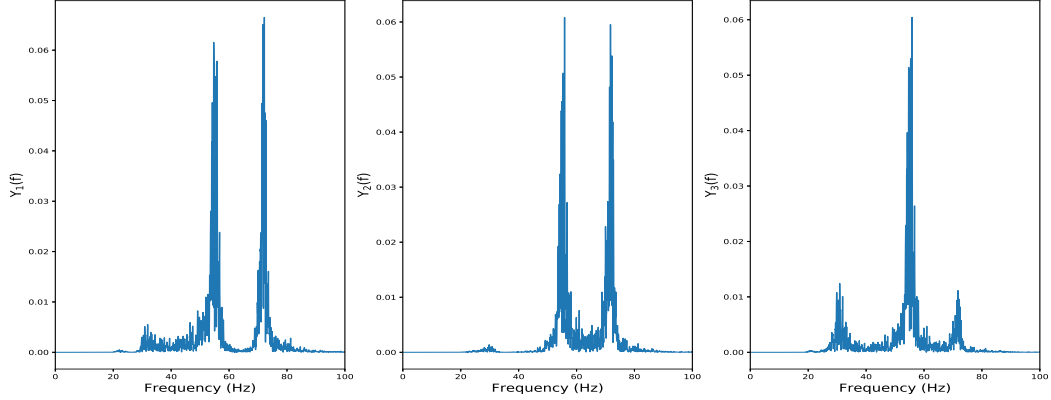


Figure 18: PSDs samples of three-floor experimental structure, physical coordinates of state 12.

Training the cycle-GAN using the accelerations from all 50 experiments, the decomposition achieved is shown in the bottom plots of Figure 19. The model that yielded the results shown had 100 nodes in its hidden layers.

510 The decoupling of the modes is clear. Each latent variable corresponds to a different mode and the algorithm performs better than using a PCA decomposition (shown in the top row of the same figure). Each latent variable's PSD has a clearly dominant peak. In the first, and especially in the second, plot of the PCA decomposition PSDs, the modes are coupled and none is dominant.

515

4.4.2. Experimental data: State 14

Following the same procedure, using the State 14 data (whose natural coordinates PSDs are shown in Figure 20), the results of the modal decomposition are shown in Figure 21. The comparison between the cycle-GAN decomposition and the PCA decomposition clearly indicate that the former has achieved

520 better results. The separation of the modes is almost perfect compared to the PCA case, where the yielded coordinates clearly do not achieve any modal separation.

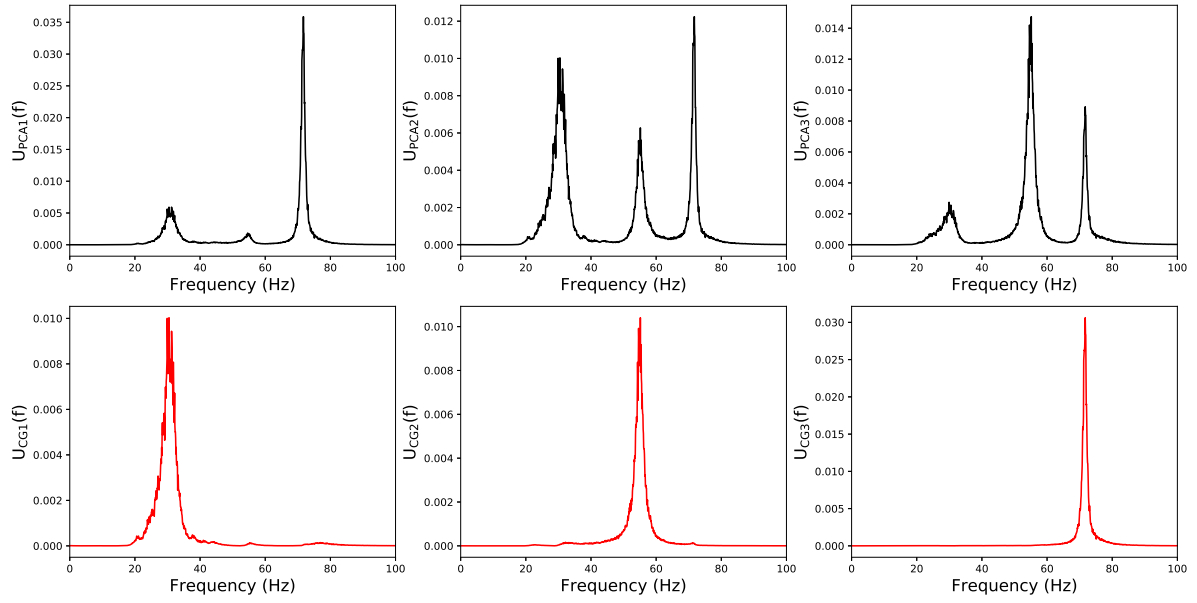


Figure 19: PSDs of three-floor experimental structure, PCA decomposition (top) and cycle GAN decomposition selected via the inner product criterion (bottom); state 12.

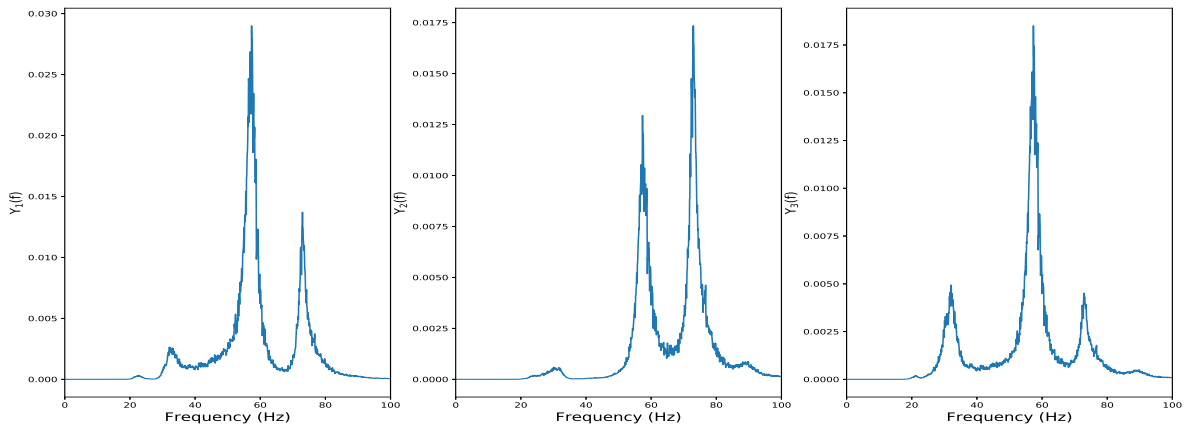


Figure 20: Average PSDs of three-floor experimental structure, physical coordinates of state 14.

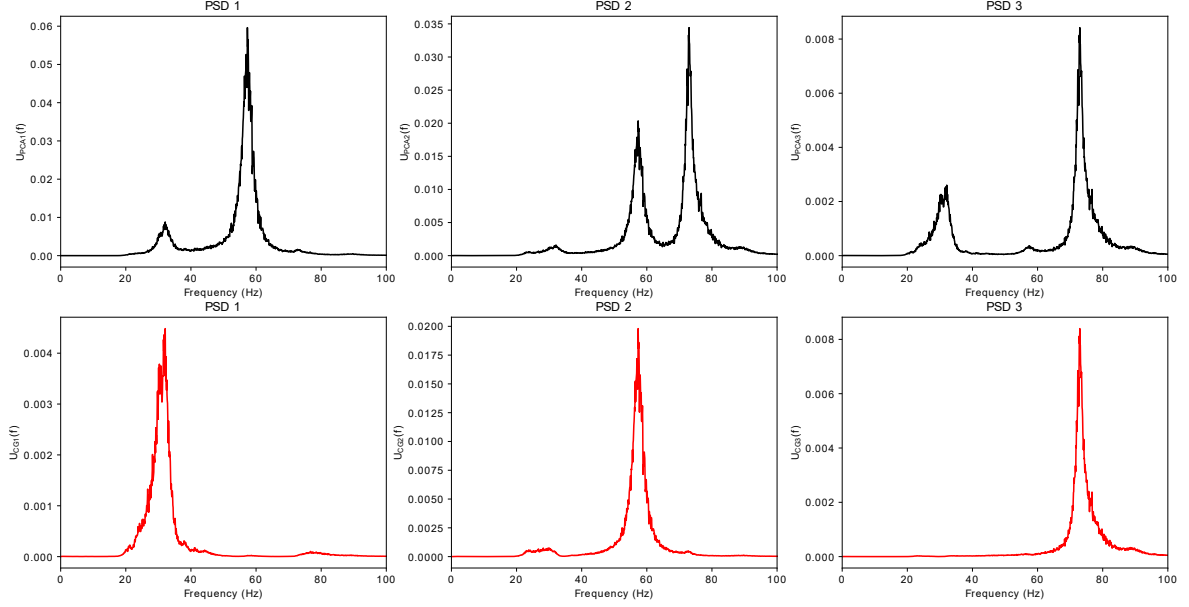


Figure 21: PSDs of three-floor experimental structure, PCA decomposition (top) and cycle GAN decomposition selected via the inner product criterion (bottom); state 14.

5. Superposition

525 The second essential part of modal analysis is the superposition step. For linear modal analysis, superposition is defined simply as the summation of the displacements that correspond to each mode. In the case of nonlinear modes, clearly it cannot be that simple. A nonlinear superposition function has to be defined in order to map coordinates from the modal space back to the physical. Definition of a superposition function, in the current work, is a
 530 major addition to the method compared to [10, 11].

Since the cycle-GAN was used, this mapping has already been defined. The inverse mapping is the goal of the second generator. The second generator has been trained in parallel with the generator that was used to decompose the movement of the structure into the modal space. Subsequently, the
 535 performance of the inverse mapping for the presented case studies is shown and discussed.

For each case study, the superposition mapping is provided by the generator $G_{U \rightarrow Y}$. The generator selected is taken from the same training epoch as the forward mapping generator $G_{Y \rightarrow U}$. To evaluate the superposition efficiency, a *normalised mean-square error* (NMSE) for each reconstruction is computed

by,

$$NMSE = \frac{100}{N\sigma_y^2} \sum_{i=1}^N (\hat{y}_i - y_i)^2 \quad (10)$$

where N is the total number of displacements samples, σ_y^2 is the variance of the displacements, y_i is the real recorded displacement and \hat{y}_i is the superposition provided by $G_{Y \rightarrow U}(G_{U \rightarrow Y}(y))$. The NMSE in each case study is computed using all the displacement samples from every degree of freedom of the systems. The NMSE is a convenient measure of error in regression problems, since it provides an objective measure of the accuracy, regardless of the scale of the data. NMSE values close to 100% indicate that the model does no better than simply using the mean value of the data, while the lower the value the better the model is calibrated. From experience, values of NMSE lower than 5% indicate a well-fitted model, and values lower than 1% show an excellent model.

5.1. Superposition for the two-degree-of-freedom system

Part of the results of superposition for the two-degree-of-freedom system are shown in Figure 22. The inverse-mapping function provided by the generator $G_{Y \rightarrow U}$ is very accurate; This is confirmed by the NMSE value, which for this case is 0.46%.

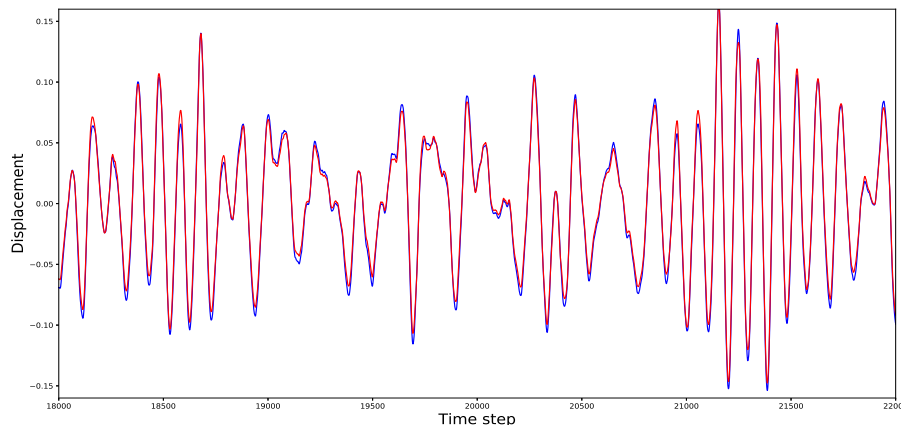


Figure 22: Superposition/inverse modal transformation for the two-degree-of-freedom system (red) and original displacements (blue).

5.2. Superposition for the three-degree-of-freedom system

For the three-degree-of-freedom system, some results are shown in Figure 23. The inverse mapping this time seems a little less accurate visually; however,

the NMSE value for this case is still only 0.294%.

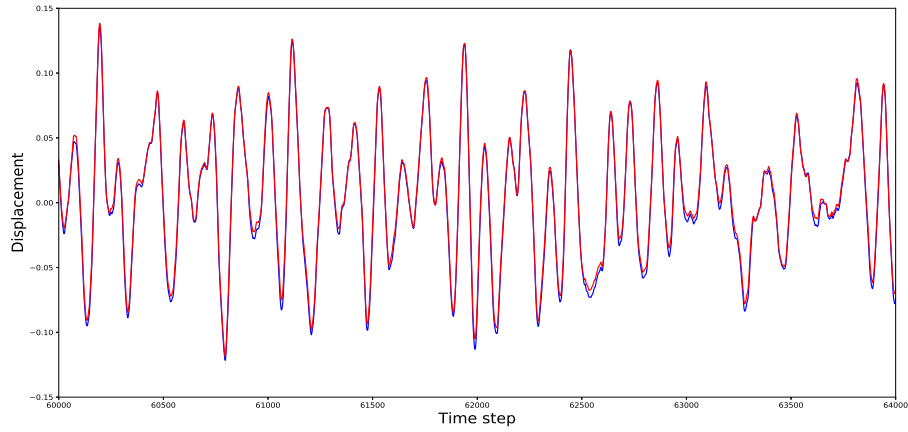


Figure 23: Superposition/inverse modal transformation for the three-degree-of-freedom system(red) and original displacements (blue).

5.3. Superposition for the four-degree-of-freedom system

For the four-degree-of-freedom system, some results are shown in Figure 24. The NMSE value for this case is 1.92%. This error is higher than the previous cases, but still implies a good inverse mapping.

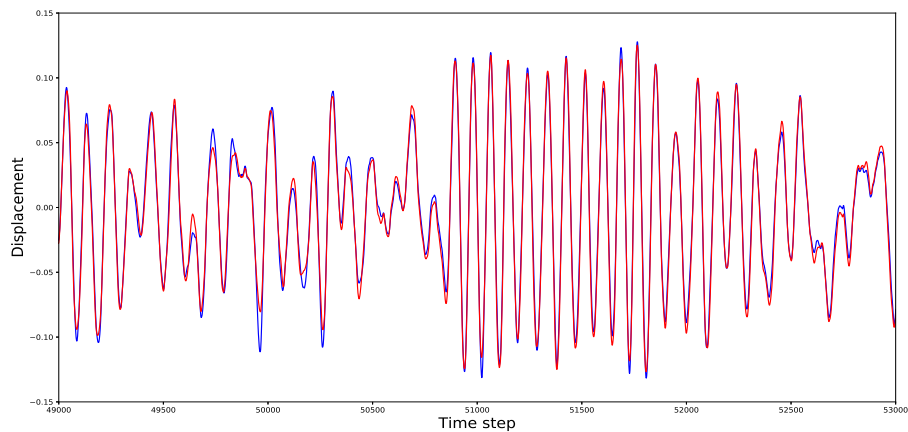


Figure 24: Superposition/inverse modal transformation for the four-degree-of-freedom system(red) and original displacements (blue).

5.4. Superposition for the experimental system

5.4.1. Experimental data: State 12

For the State 12 of the experimental system, some results are shown in Figure 25; the NMSE value for this case is 1.23%.

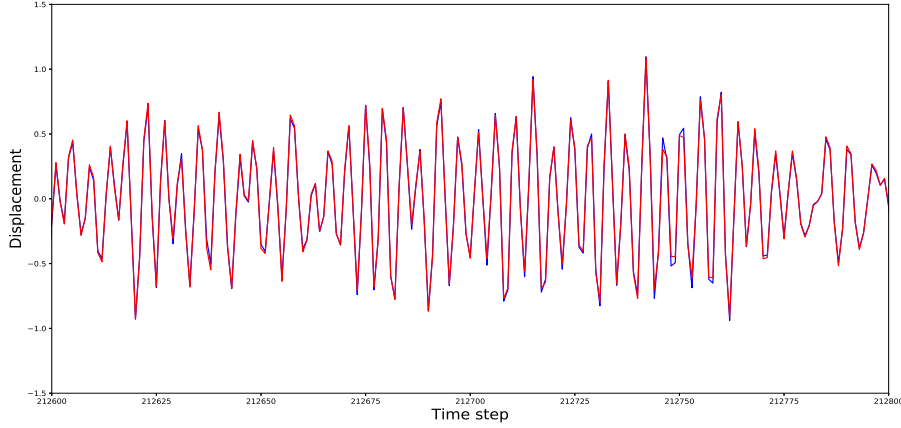


Figure 25: Superposition/inverse modal transformation for the experimental system (red) and original displacements (blue), state 12.

5.4.2. Experimental data: State 14

For the State 14 of the experimental system, some results are shown in Figure 26. The NMSE value for this case is 10.93%. Clearly the performance of the algorithm is not as good as in the previous examples. This might be explained by the harsh and highly-nonlinear nature of this final state of the experimental set up.

6. Modal correlation study

In a previous section it was mentioned that modal coordinates are expected to be independent, since the modal space is pre-defined as sampling from variables $Y \sim \mathcal{N}(\boldsymbol{\mu}, I_n)$. However, the algorithm is searching for a mapping from the modal space to the natural without any constraint regarding which samples of the modal space to use. At the same time, the algorithm has to balance three types of losses (adversarial, reconstruction and orthogonality), and may result in overlooking one of them in some degree. The adversarial loss, as discussed, enforces the statistical independence of the modal coordinates and, if overlooked in some degree, the result may be that the modal coordinates are correlated. In the current section, a further study of the correlation of the modal coordinates is performed.

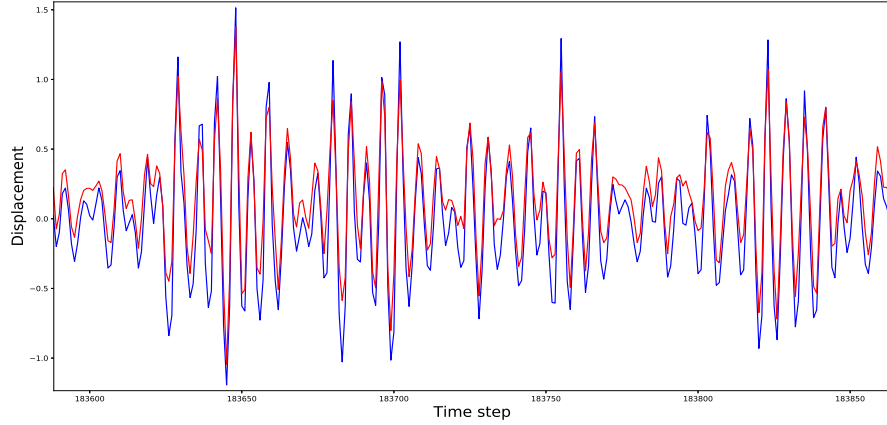


Figure 26: Superposition/inverse modal transformation for the experimental system (red) and original displacements (blue), state 14.

6.1. Linear correlation

585 The most simple case of correlation measure here will be Pearson's correlation coefficient, which for a pair of random variables (X, Y) is given by,

$$\rho_{X,Y} = \frac{\mathbb{E}[(X - \mu_X)(Y - \mu_Y)]}{\sigma_X \sigma_Y} \quad (11)$$

where μ_X, μ_Y are the mean values of X and Y respectively, $\mathbb{E}[\]$ is the expected value of the quantity in brackets and σ_X, σ_Y the standard deviations of X and Y . The correlation coefficient ρ takes values in the interval $[-1, 1]$.
 590 Values closer to -1 and 1 mean that the samples are highly correlated, while values closer to 0 mean that the two random variables are less correlated (in a linear manner).

Using such a correlation measure between observations of two variables reveals only potential linear dependency between the variables. For linear
 595 modal analysis, such a measure is minimised for the modal coordinates. Since linear modal analysis is equivalent to PCA, this is an expected effect of the transformation. Therefore, to study correlation in the case of nonlinear modal analysis, a nonlinear correlation measure should also be used.

6.2. Nonlinear correlation

The correlation measure selected here to study the nonlinear modes is *distance correlation* [25]. The distance correlation can be thought of as a generalisation of Pearson's correlation coefficient, and is a way of detecting higher-order

correlations between data. In order to calculate the distance correlation for two random variables (X, Y) , two distance matrices A and B have to be defined. To do this, elements $\alpha_{j,k}$ and $\beta_{j,k}$ are first defined by,

$$\begin{aligned}\alpha_{j,k} &= \|x_j - x_k\|, & k, j &= 1, 2, \dots, n \\ \beta_{j,k} &= \|y_j - y_k\|, & k, j &= 1, 2, \dots, n\end{aligned}\tag{12}$$

where n is the number of observations. Subsequently, the matrices A and B are computed as,

$$\begin{aligned}A_{j,k} &= \alpha_{j,k} - \bar{\alpha}_j - \bar{\alpha}_k + \bar{\alpha}.., & k, j &= 1, 2, \dots, n \\ B_{j,k} &= \beta_{j,k} - \bar{\beta}_j - \bar{\beta}_k + \bar{\beta}.., & k, j &= 1, 2, \dots, n\end{aligned}\tag{13}$$

where $\bar{\alpha}_k$ is the mean value of the k th column, $\bar{\alpha}_j$ the mean of the j th row and $\bar{\alpha}..$ the mean value of all α elements. Having computed the two matrices, the *distance covariance* is given by,

$$dCov^2(X, Y) = \frac{1}{n^2} \sum_{j=1}^n \sum_{k=1}^n A_{j,k} B_{j,k}\tag{14}$$

Defining as *distance variance* the distance covariance of a variable to itself ($dVar^2(X) = dCov^2(X, X)$) the distance correlation is calculated by,

$$dCor(X, Y) = \frac{dCov(X, Y)}{\sqrt{dVar(X)dVar(Y)}}\tag{15}$$

600 The distance correlation measure is bounded in the interval $[0, 1]$. Lower values indicate independence of the random variables while higher values indicate higher dependence. In the following subsections, both correlation metrics will be computed for the modal coordinates defined by the algorithm proposed here.

605 6.3. Correlation in the case studies

6.3.1. Correlation of the two-degree-of-freedom modes

Calculating the correlation coefficients for the modal coordinates of the two-degree-of-freedom system, the results are plotted in a ‘heat-map’ coloured matrix in Figure 27. On the diagonal, obviously, the correlation coefficients
610 are equal to 1 since they are the autocorrelation coefficients of the modal coordinates. The values of the coefficients between the two modes are quite low indicating that independence of the modes has been achieved to a satisfactory degree.

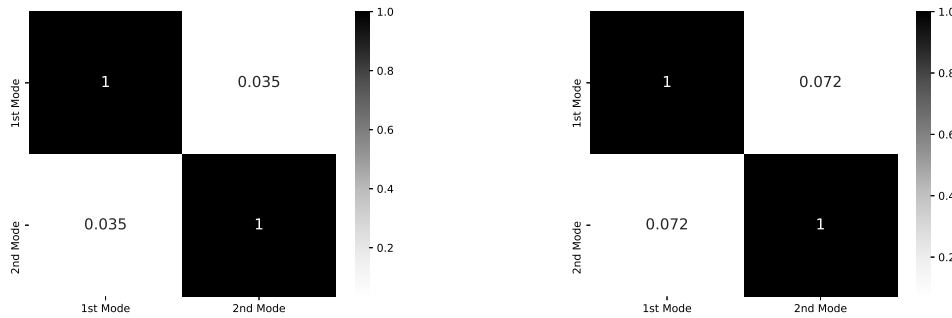


Figure 27: Pearson's linear correlation coefficient (left) and distance correlation coefficient (right) of the modal decomposition computed for the two-degree-of-freedom system.

6.3.2. Correlation of three-degree-of-freedom modes

615 For the three-degree-of-freedom system, the correlation coefficients are shown in Figure 28. In this case the correlation coefficients are also quite low.

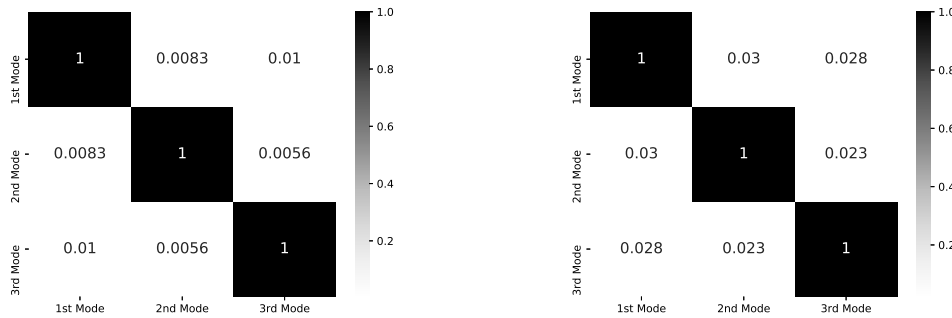


Figure 28: Pearson's linear correlation coefficient (left) and distance correlation coefficient (right) of the modal decomposition computed for the three-degree-of-freedom system.

6.3.3. Correlation of four-degree-of-freedom modes

620 For the four-degree-of-freedom system, the correlation coefficients are shown in Figure 29. The correlation coefficients for this case also appear to be low enough to assume that the modal coordinates are basically uncorrelated.

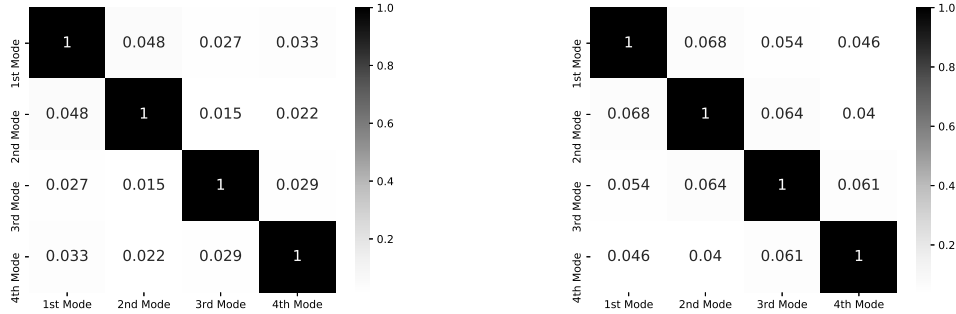


Figure 29: Pearson's linear correlation coefficient (left) and distance correlation coefficient (right) of the modal decomposition computed for the four-degree-of-freedom system.

6.3.4. Experimental data: State 12

For the State 12 of the experimental system, the correlation coefficients are shown in Figure 30. Again, the distance correlation and Pearson's correlation coefficient have low values.

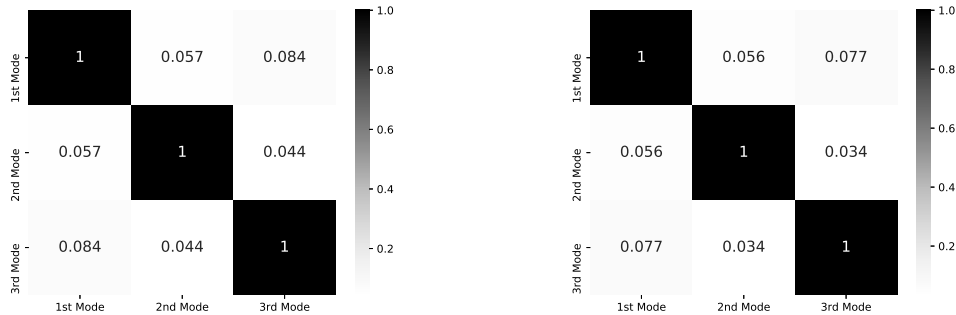


Figure 30: Pearson's linear correlation coefficient (left) and distance correlation coefficient (right) of the modal decomposition computed for the experimental system.

625 6.3.5. Experimental data: State 14

For the State 14 of the experimental system, the correlation coefficients are shown in Figure 31. This time the correlation values are higher than before; this might also be the result of the highly-nonlinear system. The highest correlation is observed between the first and the second modal coordinates.

630 It can also be explained by energy observed in the PSD of the second modal coordinate (Figure 21) close to the frequencies of the first mode.

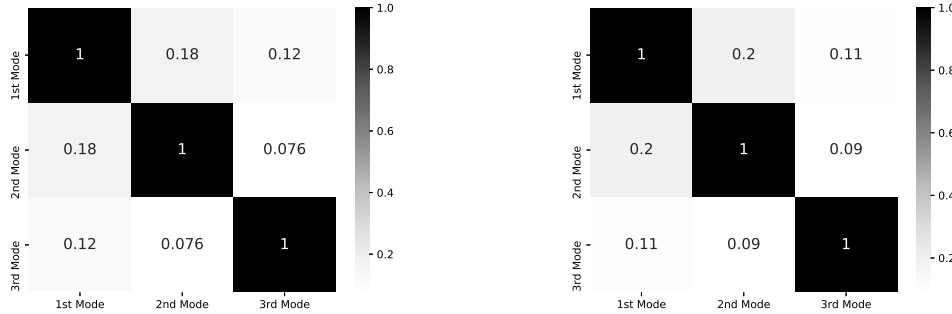


Figure 31: Pearson's linear correlation coefficient (left) and distance correlation coefficient (right) of the modal decomposition computed for the experimental system.

7. Conclusions

An algorithm was described as an alternative of the framework developed in [10]. The algorithm aims at performing data-driven nonlinear modal analysis.

635 The concept on which the nonlinear modal analysis is based is that of decomposition into SDOF systems. However, instead of decomposing the target dataset of displacements (or accelerations) into a modal space, a mapping from a pre-defined modal space onto the dataset is attempted. Both forward and backward mappings are learnt using a cycle-GAN model. In contrast to
 640 the classic application of such a model, a restriction of orthogonality is also imposed here. This restriction corresponds to the orthogonality of the mode shapes. Moreover, a model selection criterion different to the loss function used during training, based on the inner product of two PSDs, is used to pick the model that separates the modes most efficiently.

645 The selection criterion, apart from selecting the model that best separates the modes, is also expected to implicitly minimise the correlation of the modes. Given that a modal variable reacts only to some interval of frequencies that is not overlapping with the rest of the modal coordinates, its statistical correlation with the rest of the modal variables is minimised. The
 650 minimisation of the correlation is also implicitly imposed via the adversarial training. The target distribution of the generator, which transforms from natural coordinates to modal, is a multivariate Gaussian with diagonal correlation matrix. The adversarial training is forcing the generator to generate samples with similar distribution and therefore with correlation between
 655 modal coordinates close to zero.

Via three simulated case studies and two experimental, the performance of the algorithm was illustrated. The proposed model was able to decouple modes in two, three and four degree-of-freedom systems with cubic stiffening nonlinearities and also in two bilinear (bumper column) experimental cases.

660 The main reason that the algorithm performs better than the algorithm
proposed in [10] is probably that there are no restrictions on the polynomial
order of the composition (or decomposition) function used or on the objective
function.

Parallel to the training of the decomposition, the superposition function is
665 also trained and as illustrated, yields accurate results (below 2% normalised
mean square error, except for the second experiment, implying very good
accuracy models); thus providing a solution to invertibility problems in [10].
The second generator can be used to map from the modal space to the natural
coordinate space. The only case that a higher error in the superposition
670 function was observed was the most nonlinear experimental case. The
nonlinearity in the latter case is imposed by the bumper and the column
defining a bilinear nonlinearity. When the bumper is too close to the column
and too close to the equilibrium point, the nonlinearity is also imposed by
the striking of the two elements, which causes sudden energy dissipation
675 and non-proportional damping, as well as additional vibration components
to the structure. Nonlinearity on non-proportional damping might require
considering also the velocity variables within the algorithm, in order to
achieve a more efficient decomposition. The problem might also be that the
algorithm did not have enough data to satisfy all the constraints imposed
680 during training (adversarial loss, orthogonality loss and reconstruction loss)
given how highly and non-smoothly (sudden stiffness change and energy
dissipation) nonlinear the case is. In any case, if the accuracy is not sufficient,
one can train a separate regression algorithm (for example a Gaussian process
or a regression neural network), in order to define a superposition function.

685 Two correlation coefficients (Pearson’s correlation coefficient and the distance
correlation) are calculated for the modal coordinates that were computed
by the algorithm. The values of the coefficients seem to confirm the initial
assumption, that the *a priori* selection of a modal space and the application
of PCA on the physical coordinates forces the modal coordinates to be
690 uncorrelated. Regarding the correlations, again in the most nonlinear case
higher values of the correlation metrics were observed. The high correlation
might be the result of lack of data as well as the largely nonlinear system
in combination with a non-smooth nonlinearity. The distance correlation is
used because Pearson’s correlation can be ‘fooled’ by nonlinear correlations.
695 The distance correlation cannot be tricked by higher-order correlation and
this is verified by the values on the off-diagonals that are generally higher
than the corresponding Pearson’s correlation.

In total, the proposed algorithm solves most of the problems that remained in
[10]. The algorithm is able to perform for higher-degree-of-freedom systems
700 than the one presented in [10] and also for the experimental data. Moreover,
the dot product criterion presented replaces the ‘by eye’ selection of the model

that best separates the modes. The problem of superposition is also solved, since by using the cycle-GAN algorithm, the superposition function is trained in parallel. Finally, the modal coordinates seem to have low correlation values
705 except for the harshly-nonlinear case, an issue that remains to be solved.

Acknowledgements

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 764547. KW would like to thank the UK Engineering and
710 Physical Sciences Research Council (EPSRC) for an Established Career Fellowship (EP/R003645/1). DW would like to acknowledge the support of EPSRC grand EP/R006768/1.

References

- [1] D.J. Ewins. *Modal Testing: Theory, Practice and Application*. John
715 Wiley & Sons, 2009.
- [2] Anil K. Chopra. *Dynamics of Structures*. Pearson Education India, 2007.
- [3] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 2(1-3):37–52, 1987.
- [4] A. Chatterjee. An introduction to the proper orthogonal decomposition. *Current Science*, pages 808–817, 2000.
720
- [5] F. Poncelet, G. Kerschen, J.-C. Golinval, and D. Verhelst. Output-only modal analysis using blind source separation techniques. *Mechanical Systems and Signal Processing*, 21(6):2335–2358, 2007.
- [6] I. Kovacic and M.J. Brennan. *The Duffing Equation: Nonlinear Oscillators and Their Behaviour*. John Wiley & Sons, 2011.
725
- [7] K. Worden and G.R. Tomlinson. *Nonlinearity in Structural Dynamics*. Institute of Physics Press, 2001.
- [8] A. Vakakis. Non-linear normal modes (NNMs) and their applications in vibration theory: an overview. *Mechanical Systems and Signal Processing*, (11):3–22, 1997.
730
- [9] Y.V. Mikhlin and K.V. Avramov. Nonlinear normal modes for vibrating mechanical systems. review of theoretical developments. *Applied Mechanics Reviews*, 63(6), 2010.

- 735 [10] K. Worden and P.L. Green. A machine learning approach to nonlinear modal analysis. *Mechanical Systems and Signal Processing*, 84:34–53, 2017.
- [11] N. Dervilis, E. Thomas T. Simpson, D.J. Wagg, and K. Worden. Non-linear modal analysis via non-parametric machine learning tools. *Strain*, 740 55(1):e12297, 2019.
- [12] S.W. Shaw and C. Pierre. Normal modes for non-linear vibratory systems. *Journal of Sound and Vibration*, 164(1):85–124, 1993.
- [13] J.Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. 745
- [14] V. A. Epanechnikov. Non-parametric estimation of a multivariate probability density. *Theory of Probability & Its Applications*, 14(1):153–158, January 1969.
- 750 [15] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S.Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- [16] Charles R Farrar and Keith Worden. *Structural Health Monitoring.: A Machine Learning Perspective*. John Wiley & Sons, 2012. 755
- [17] M. A. Kramer. Nonlinear principal component analysis using autoassociative neural networks. *AIChE Journal*, 37(2):233–243, 1991.
- [18] J. Müller, R. Klein, and M. Weinmann. Orthogonal Wasserstein GANs. *arXiv preprint arXiv:1911.13060*, 2019.
- 760 [19] W. Li, L. Fan, Z. Wang, C. Ma, and X. Cui. Tackling mode collapse in multi-generator GANs with orthogonal vectors. *Pattern Recognition*, 110:107646.
- [20] B. C. Csáji et al. Approximation with artificial neural networks. *Faculty of Sciences, Eötvös Loránd University, Hungary*, 24(48):7, 2001.
- 765 [21] P. Welch. The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, 15(2):70–73, 1967.
- [22] L. Tarassenko. *Guide to Neural Computing Applications*. Elsevier, 1998.

- 770 [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization.
arXiv preprint arXiv:1412.6980, 2014.
- [24] E. Figueiredo, G. Park Gyuhae, J. Figueiras, C. Farrar, and K. Worden
Keith. Structural health monitoring algorithm comparisons using stan-
775 dard data sets. Technical report, Los Alamos National Lab.(LANL),
Los Alamos, NM (United States), 2009.
- [25] G. J. Székely and M.L. Rizzo. The distance correlation t-test of indepen-
dence in high dimension. *Journal of Multivariate Analysis*, 117:193–213,
2013.