

Cronfa - Swansea University Open Access Repository

This is an author produced version of a paper published in:

Nonlinear Dynamics

Cronfa URL for this paper:

<http://cronfa.swan.ac.uk/Record/cronfa51462>

Paper:

Renson, L., Sieber, J., Barton, D., Shaw, A. & Neild, S. (2019). Numerical continuation in nonlinear experiments using local Gaussian process regression. *Nonlinear Dynamics*

<http://dx.doi.org/10.1007/s11071-019-05118-y>

This item is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Copies of full text items may be used or reproduced in any format or medium, without prior permission for personal research or study, educational or non-commercial purposes only. The copyright for any work remains with the original author unless otherwise specified. The full-text must not be sold in any format or medium without the formal permission of the copyright holder.

Permission for multiple reproductions should be obtained from the original author.

Authors are personally responsible for adhering to copyright and publisher restrictions when uploading content to the repository.

<http://www.swansea.ac.uk/library/researchsupport/ris-support/>

Numerical Continuation in Nonlinear Experiments using Local Gaussian Process Regression

L. Renson¹ · J. Sieber² · D.A.W. Barton¹ · A.D. Shaw³ · S.A. Neild¹ ·

Received: date / Accepted: date

Abstract Control-based continuation (CBC) is a general and systematic method to probe the dynamics of nonlinear experiments. In this paper, CBC is combined with a novel continuation algorithm that is robust to experimental noise and enables the tracking of geometric features of the response surface such as folds. The method uses Gaussian process regression to create a local model of the response surface on which standard numerical continuation algorithms can be applied. The local model evolves as continuation explores the experimental parameter space, exploiting previously captured data to actively select the next data points to collect such that they maximise the potential information gain about the feature of interest. The method is demonstrated experimentally on a nonlinear structure featuring harmonically-coupled modes. Fold points present in the response surface of the system are followed and reveal the presence of an isola, i.e. a branch of periodic responses detached from the main resonance peak.

Keywords nonlinear experiment · control-based continuation · regression-based continuation · Gaussian process regression · active data selection

1 Introduction

Numerical continuation is a popular and well-established method to systematically investigate the behaviour of

¹ L. Renson (✉), D.A.W. Barton and S.A. Neild
Faculty of Engineering, University of Bristol, Bristol, UK.
E-mail: l.rendon@bristol.ac.uk

² J. Sieber
Centre for Systems, Dynamics and Control, College of Engineering, University of Exeter, UK.

³ A.D. Shaw
College of Engineering, Swansea University, UK.

nonlinear dynamical systems and perform bifurcation analysis [1, 2]. At a basic level, numerical continuation finds the solutions of a zero-problem $\mathbf{f}(\mathbf{x}, \lambda) = \mathbf{0}$ where \mathbf{x} are the system states, and tracks the evolution of the solutions as the parameter λ is varied. Based on a mathematical model, the long-term behaviours of a system, such as steady-states and periodic responses, can easily be represented by such a zero-problem. Bifurcations, and hence stability changes, in those behaviours can then be detected along the solution path and in turn tracked by adding suitable constraint condition(s) and free parameter(s). The principles of numerical continuation are extremely general such that the method has been applied to a wide range of problems across engineering and the applied sciences as, for instance, in bio-chemistry [3], physics [4], mechanics [5] and fluid dynamics [6].

Without the need for a mathematical model, control-based continuation (CBC) is a means to define a zero-problem based on the inputs and outputs of an experiment, thereby allowing the principles of numerical continuation to be applied to a physical system directly during experimental tests. The fundamental idea of CBC is to use feedback control to stabilise the dynamics of the experiment whilst making the control system *non-invasive* such that it does not modify the position in parameter space of the responses of the open-loop experiment of interest. This non-invasiveness requirement defines a zero-problem whose solutions can be found and tracked in the experiment using the same path-following principles and methods as in the numerical context.

CBC is similar in principle to other methods such as the famous OGY (Ott, Grebogi, Yorke) control tech-

nique that has been extensively used to stabilise unstable periodic responses embedded in chaotic attractors [7]. The OGY method was coupled to continuation algorithms in [8], but the application of this method to general nonlinear experiments¹ remains challenging due to the particular form of control used in the OGY technique. Other examples of control techniques used to experimentally measure unstable responses are the Pyragas delayed feedback control [9] and phase-locked loops [10]. The latter has recently been applied to several nonlinear mechanical systems [11, 12, 13]. Although these methods and CBC share a number of similarities, CBC does not assume any particular form of control.

CBC was first proposed by Sieber and Krauskopf [14], and experimentally demonstrated on a parametrically-excited pendulum [15]. The method has since been successfully applied to a range of mechanical systems, including an impact oscillator [16, 17, 18], oscillators with magnetic nonlinearities [19, 20] and a cantilever beam with a nonlinear mechanism at its free tip [21]. Through those studies, CBC proved to be a versatile technique that can extract important nonlinear dynamic features such as nonlinear frequency response curves [22], nonlinear normal modes [23, 24] and isola [21] directly in the experiment. However, the systematic application of CBC to general nonlinear experiments remains challenging. Most existing continuation algorithms are ideal only in a numerical context where the solution path is smooth and derivatives can be evaluated to high precision. This is not easily achievable in experiments where solutions and derivative estimates are corrupted by measurement noise. Schilder *et al.* [25] discussed the effect of noise on continuation algorithms. In particular, the tangential prediction and orthogonal correction steps of the commonly-used pseudo-arc length continuation algorithm were shown to perform poorly in a noisy experimental context. Similarly, continuation step size control techniques that reduce continuation steps when convergence is not achieved were also shown to be inadequate as reducing the step size usually makes noise distortions even more apparent. Schilder *et al.* proposed alternative numerical strategies that are more robust to noise [25]. These new strategies significantly improve the robustness of CBC as they are able to trace solution paths that are no longer smooth due to noise. These strategies work well for low levels of noise but are not sufficiently robust to find and track dynamic features that are very sensitive to noise as, for instance, bifur-

cations which are defined in terms of derivatives.

This paper proposes an algorithm that not only makes CBC more robust to noise but also enables the tracking of general dynamic features that are not directly measurable in the experiment due to perturbations from noise for instance. The proposed approach is fundamentally different from the one taken by Schilder and co-workers. Multivariate regression techniques are exploited to locally model the response surface of the experiment, providing local models that are smooth and cheap to evaluate and thus enabling the use of established numerical continuation techniques. The proposed method is inspired by the work reported in [20] where a single-parameter cubic polynomial regression was exploited to capture the geometry of the response surface at a fixed forcing frequency and used to detect and track a limit-point bifurcation in parameter space. The algorithm presented in [20] is however limited to bifurcation curves with simple geometries. For instance, the algorithm cannot follow the bifurcation through a cusp. The method proposed here overcomes this limitation.

A key difficulty in using online models is to devise a strategy to collect the experimental data necessary to build them. We use Gaussian Process regression (GPR) techniques to address this challenge. GPR has many desirable features, such as the ease of extension to models with multiple inputs and outputs, the ease of expressing uncertainty, the ability to capture a wide range of behaviours using a simple (hyper-)parameterisation, and a natural Bayesian interpretation. Here, based on the data points already captured, GPR will allow us to determine where to collect new data points to maximise the potential information they will provide about the dynamic feature of interest. This active selection of the data based on our current knowledge and the feature of interest contrasts with the approaches currently found in the literature [26, 27] where data collection and identification are two activities often performed separately.

The proposed algorithm is presented in Section 2 and demonstrated on a nonlinear mechanical structure composed of a cantilever beam with a nonlinearity attached at its free end (Sections 3 and 4). The natural frequencies of the first two modes of the structure are almost in a 3:1 ratio, which leads to strong harmonic couplings between these modes and the presence of complicated nonlinear dynamic behaviours. In particular, a branch of stable periodic responses detached from the main resonance peak can be observed. Our new algorithm is employed to track fold points present in the response surface of the system. The obtained

¹ The term “*nonlinear experiment*” refers to an experiment for which any model describing its behaviour has to be nonlinear to be consistent with observations.

curves are found to have a complex geometry due to the presence of the modal interaction and to reveal the presence of the isola.

2 Tracking dynamic features using an online regression-based continuation algorithm

The zero-problem defining the responses tracked in the experiment is now explained (Section 2.1). The algorithm used to experimentally solve and continue the solution of this zero-problem is then presented in Section 2.2. Detailed discussions of the important components of this algorithm are given in Sections 2.3 – 2.5.

2.1 Definition of the zero-problem

The response surface of an experiment such as the one considered here presents the generic form shown in Figure 1, where A is the system response amplitude and (λ_1, λ_2) are parameters. Such response surface can be easily obtained experimentally using an established CBC algorithm [20]. In this paper, we are interested in directly tracking geometric features of that response surface during the experiment and, in particular, the fold curve represented in solid black (—). Responses that lie on that curve satisfy the scalar constraint equation

$$\frac{d\lambda_2(\lambda_1, A)}{dA} = 0. \quad (1)$$

The response amplitude A is a natural choice for the parameterisation of the fold curve as it is sufficient to uniquely define it (at least in Figure 1). The parameter λ_1 serves as the free parameter for the continuation and, in the present experimental context, corresponds to the forcing frequency ω . λ_2 is the external harmonic force excitation amplitude Γ .

The choice of ω and A as independent variables in Eq. (1) also stems from the nature of these variables in the experiment. More precisely, the response amplitude A is indirectly imposed by the amplitude of the reference signal of the control system used in CBC (in particular, one of its fundamental harmonic components). As such, A and ω can both be viewed as ‘controllable inputs’ to the experiment. In contrast to continuation applied to a numerical system, the force amplitude Γ is here considered as a quantity that is difficult to set directly and hence one that is measured from the experiment rather than imposed.

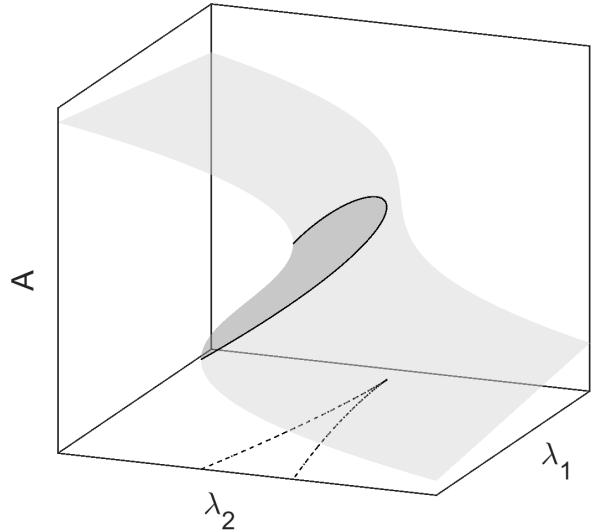


Fig. 1 Typical response surface that can be found in the experiment. A is a measure of the response of the system and λ_1 and λ_2 are parameters, here selected as $\lambda_1 = \omega$ and $\lambda_2 = \Gamma$. Fold curve (—) and its projection (---) in parameter space.

2.2 Overview of the online regression-based algorithm

Finding data points experimentally that satisfy Eq. (1) can be difficult and error prone due to the presence of noise affecting derivative calculations. This issue is addressed here by creating, *online* (i.e. while the experiment is running), regression models that capture the local dependence of the applied force amplitude as a function of the forcing frequency and response amplitude. Derivatives can be effectively and accurately calculated for these models, which in turn allows us to find and then track the fold curve in the experiment using standard numerical continuation algorithms.

The principal steps of the proposed algorithm are shown in Figure 2. The algorithm is initialised by collecting a user-defined number of data points n_0 distributed in a regular pattern around a starting point which is assumed to be close to a fold. This first set of experimental data is then used to create the first GPR model and estimate the regression hyper-parameters (Section 2.3). Using this model, a fold point is also found by solving Eq. (1) using a standard Newton-like algorithm. From this first point, standard predictor-corrector continuation algorithms can be exploited (Section 2.4). However, unlike in numerical simulations, the correction step is followed by a data collection step that aims to make the solution of the zero-problem (1) robust to new data points (Section 2.5). The addition of each new data point to the local GPR model is followed by a correction procedure to update the solution

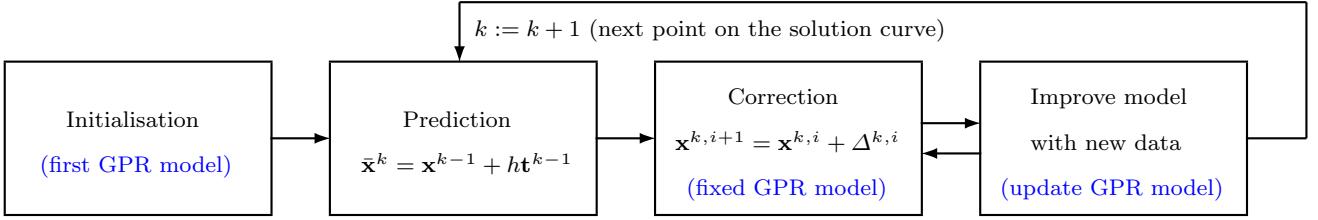


Fig. 2 Overall structure of the regression-based continuation algorithm proposed here. \mathbf{x}^k is the k^{th} solution found using the online regression-based continuation, h is the continuation step size, \mathbf{t}^k is the tangent to the solution curve at \mathbf{x}^k . $\Delta^{k,i}$ is the correction applied to $\mathbf{x}^{k,i}$ at the i^{th} iteration of the correction algorithm.

of Eq. (1). When no additional data is needed, a new prediction step is performed to progress along the solution curve.

2.3 Gaussian Process Regression

The amplitude of excitation Γ is locally modelled as a function of the response amplitude and excitation frequency using GPR. A Gaussian process is a probabilistic model that can be used to capture a wide range of nonlinear functions from input-output data without any explicit assumptions on their mathematical relationship [28]. The amplitude of excitation, Γ , is modeled by the distribution²

$$\Gamma(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), \kappa(\mathbf{x}, \mathbf{x}')), \quad (2)$$

where $\mathbf{x} = (\omega, A)$ is the vector of inputs, $m(\mathbf{x})$ and $\kappa(\mathbf{x}, \mathbf{x})$ are the mean and covariance functions. For the data point i , the measured excitation amplitude, denoted F_i , is assumed to differ from the function values Γ_i by the presence of an additive Gaussian noise with zero mean and variance σ_n^2 such that

$$F_i = \Gamma_i(\mathbf{x}) + \epsilon_i \quad \text{with} \quad \epsilon_i \sim \mathcal{N}(0, \sigma_n^2). \quad (3)$$

This ‘noise’ captures in fact measurement errors, which not only include measurement noise but also other inaccuracies that arise in the collection of data points using CBC (see Section 4).

Given a data set \mathcal{D} with n measurements

$$\{(\mathbf{x}_1, F_1), (\mathbf{x}_2, F_2), \dots, (\mathbf{x}_n, F_n)\},$$

grouped in an input matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$ and an output vector $\mathbf{F} = [F_1, F_2, \dots, F_n]^T$, the prediction

² Rigorously, our notations should distinguish the model output from the exact force amplitude Γ as the model is only an approximation of the truth. However, in the present experimental context, the exact force amplitude is unavailable to us. As such, to keep our notations simple, the model output will also be denoted Γ .

of the force amplitude at n_* unmeasured inputs \mathbf{X}_* is given by the mean of the predictive distribution as

$$\mathbf{\Gamma}_* = \boldsymbol{\kappa}(\mathbf{X}, \mathbf{X}_*)^T \boldsymbol{\kappa}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{F}. \quad (4)$$

where $(\cdot)^T$ represents the transpose operation. $\boldsymbol{\kappa}(\mathbf{X}, \mathbf{X}_*)$ corresponds to the $n \times n_*$ matrix resulting from the application of the covariance function $\kappa(\mathbf{x}_i, \mathbf{x}_{*j})$ for all i and j . Similarly for $\boldsymbol{\kappa}(\mathbf{X}, \mathbf{X})$. Note that GPR differs from parameter estimation techniques and the regression method used in [20] because measured data points are needed to perform predictions (see Eq. (4)).

An effective computation of Eq. (4) can be achieved using the Cholesky factorization to decompose the covariance matrix $\boldsymbol{\kappa}(\mathbf{X}, \mathbf{X})$ into a lower-triangular matrix and its conjugate transpose [28]. The former is then stored and used to efficiently compute the inverse for different unmeasured inputs \mathbf{x}_* . The Cholesky decomposition can also be efficiently updated when data points are added or removed from the data set \mathcal{D} as in Section 2.5. Note that the inputs (ω, A) of the GPR model are assumed to be noise-free but they are in reality measured quantities corrupted by noise. It is possible to extend the GPR model used here to address noisy inputs [29]; however this is not considered here as it was unnecessary for the system considered.

Although GPR is a non-parametric approach that does not make assumptions on the functional form of the modelled function, assumptions regarding its smoothness are introduced in the covariance function and the choice of hyper-parameters $\boldsymbol{\theta}$. The covariance function considered here is the widely-used, infinitely-differentiable squared-exponential (SE) or radial basis function (see, for instance, Eq. (5.1) in [28]). With two-dimensional inputs and observation noise, this covariance function includes four hyper-parameters $\boldsymbol{\theta} = (\sigma_n^2, \sigma_f^2, l_\omega, l_A)$. These parameters have a clear physical interpretation. In particular, l_ω and l_A represent characteristic length scales that express the distance to be travelled along a particular axis in the input space for covariance values to become uncorrelated. σ_n^2 is the variance associated with

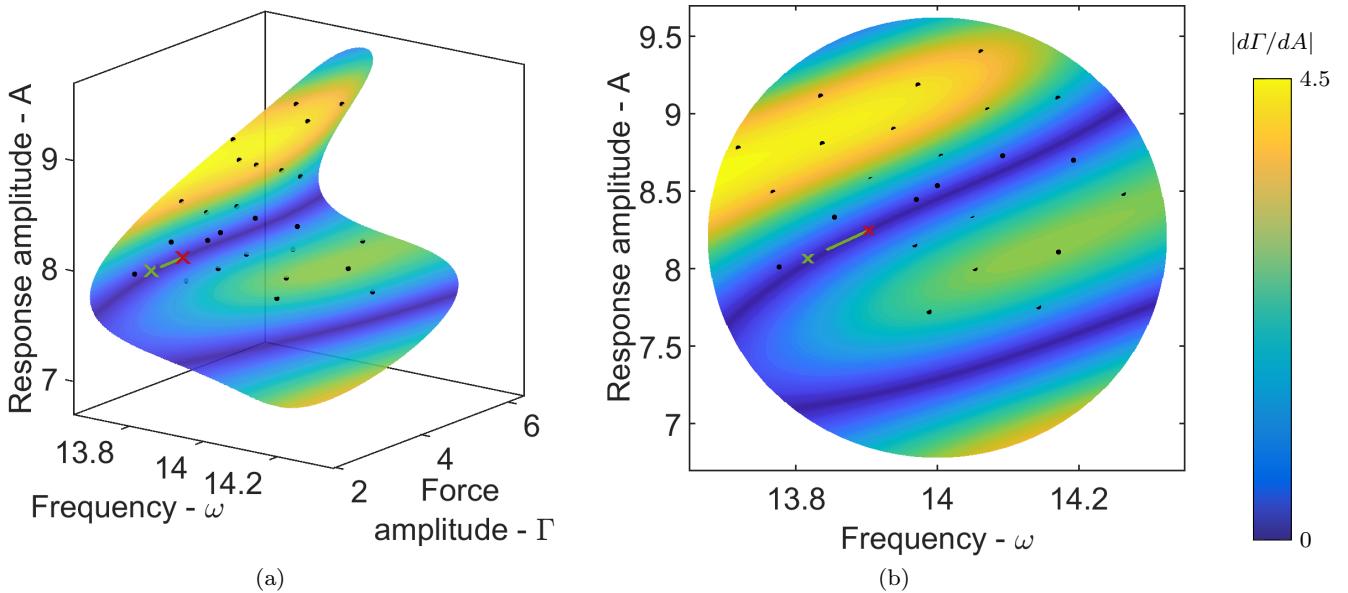


Fig. 3 Illustration of the local model obtained using GPR at the initialisation of the algorithm. (a) and (b) correspond to different views of the same model. The colour map represents the absolute value of Eq. (1), i.e. $|d\Gamma/dA|$. (●) Data points used to create the model. (✗) Fold point. (✗) Predicted fold point. (—) Tangent vector to the solution curve at fold point (✗).

measurement errors, and σ_f^2 is a magnitude factor. The hyper-parameters $\boldsymbol{\theta}$ of the models were obtained by maximizing the log marginal likelihood

$$\log p(\mathbf{F}|\mathbf{X}, \boldsymbol{\theta}) = \frac{1}{2} \mathbf{F}^T \boldsymbol{\kappa}(\mathbf{X}, \mathbf{X})^{-1} \mathbf{F} - \frac{1}{2} \log |\boldsymbol{\kappa}(\mathbf{X}, \mathbf{X})| - \frac{n}{2} \log 2\pi \quad (5)$$

where the first term represents the data fit of the model and the second term penalizes the complexity of the model [28]. The optimization was performed using a Quasi-Newton algorithm. Hyper-parameters were determined at the start of the algorithm using the first set of n_0 data points captured. They were then kept constant during a continuation run. We note that Gaussian priors on the hyper-parameters (hyper-priors) can be used. However, with limited knowledge of the actual values and hence large covariances they had limited influence on the optimal $\boldsymbol{\theta}$ values.

As an example of a GPR model, Figure 3(a) shows a local model obtained at the start of a continuation run after collecting 25 data points (●) regularly distributed in input space (ω, A) around the estimated location of a fold point. A projection of the model in the two-dimensional input space is also given in Figure 3(b). The surface is coloured according to $|d\Gamma/dA|$ to highlight the regions where Eq. (1) is satisfied. According to the model, fold points are expected in two distinct regions (in dark-blue). The one where the response amplitude is higher is where the actual folds are

located. The lower one is in fact located outside the data set and is an artifact feature created by the regression. These artifacts do not affect the algorithm as long as the continuation steps are small enough to stay within the available data sets. In Section 2.5, we will discuss how to improve such GPR models with additional experimental data but first we discuss the continuation approach.

2.4 Numerical continuation

The continuation problem is to solve and track the solutions of Eq. (1). Starting from a known solution \mathbf{x}^{k-1} , the next point along the solution path is predicted to be

$$\tilde{\mathbf{x}}^k = \mathbf{x}^{k-1} + h \mathbf{t}^{k-1} \quad (6)$$

where h is the continuation step size and \mathbf{t}^{k-1} is the tangent vector to the solution curve at \mathbf{x}^{k-1} . The prediction will not in general satisfy the zero-problem such that the prediction must be corrected using a Newton-like algorithm. However, to apply Newton iterations, an extra equation has to be added to Eq. (1). The equation used here is the so-called pseudo-arc length condition

$$\mathbf{t}^{k-1} \cdot (\mathbf{x} - \mathbf{x}^{k-1}) - h = 0, \quad (7)$$

which constraints the corrections made to $\tilde{\mathbf{x}}^k$ to be perpendicular to \mathbf{t}^{k-1} .

Figure 3 illustrates a fold point (\times) predicted from a previously computed point (\times). The tangent vector used for this prediction is shown in green ($\textcolor{green}{-}$). It was found that GPR models could diverge quickly from the actual response surface outside the set of collected data points (see Section 2.3). As such, the size of the continuation steps were taken such that the prediction and correction iterations remained within the cloud of available data points. Note that it is possible to allow larger continuation steps that leave the current data set by collecting new data points after the prediction step, although this is not used here. After finding a solution to Eqs. (1) and (7), additional data points are collected to refine this solution — this is now discussed.

2.5 Improve solution with new data

To make sure that the solution of Eq. (1) is not an artifact of the model and does not critically depend on the current data, new data points are collected after the correction step of the continuation algorithm. To this end, a fixed number n_{test} of prospective data points, \mathbf{x}_*^i with $i = 1, \dots, n_{test}$, uniformly distributed in an ellipse around the current solution of the continuation problem are considered for data collection. The principal axis of this elliptical domain were chosen equal to twice the length scale hyper-parameters l_ω and l_A . The data point that most influences the results is deemed to be the most interesting point and is then experimentally collected using the established CBC technique summarised in Section 4.1.

To determine the sensitivity of the zero-problem (1) to new experimental data, an artificial measurement of the force amplitude, $\bar{\mathbf{F}}_i$, is created for each candidate data point \mathbf{x}_*^i using

$$\bar{\mathbf{F}}^i = \Gamma_*^i(\mathbf{x}_*^i) + \sqrt{\text{var}[\Gamma_*^i(\mathbf{x}_*^i)]}, \quad (8)$$

where Γ_*^i depends on currently available data and is given by Eq. (4). The variance at a particular input point \mathbf{x}_*^i is given by

$$\text{var}[\Gamma_*^i] = \kappa(\mathbf{x}_*^i, \mathbf{x}_*^i) - \kappa(\mathbf{X}, \mathbf{x}_*^i)^T \kappa(\mathbf{X}, \mathbf{X})^{-1} \kappa(\mathbf{X}, \mathbf{x}_*^i), \quad (9)$$

which is independent of any previous and future measurement of the force amplitude. The artificial measurements are individually added to the GPR model and their effect on the zero-problem assessed using

$$\beta = \left| \frac{d\tilde{\Gamma}}{dA}(\mathbf{x}^k) - \frac{d\Gamma}{dA}(\mathbf{x}^k) \right| \quad (10)$$

where \mathbf{x}^k is the current solution of the continuation problem, $\tilde{\Gamma}$ is the GPR model including the artificial measurement and Γ is the GPR model without (i.e. containing only the experimental data). β is the zero-problem sensitivity to new data and is directly used to assess the potential information gained by each artificial measurement.

Figure 4 illustrates this data selection approach. The colourmap represents the sensitivity, β , of the zero-problem (1) to a new data point. Starting from Figure 4(a), a new data point is collected where Eq. (10) is the largest (\times). This new data point is added to the GPR model and the correction step of the continuation algorithm repeated. Once a solution is found, the influence of a new data point on the new GPR model is again assessed (Figure 4(b)). The region where the first additional data point was added is now observed to be significantly less influenced by any new point. According to the model predictions, informative data points are now located in another region where a second data point is eventually recorded. Following the same procedure, data points are collected until Eq. (10) is below a user-defined tolerance across the whole region as in Figure 4(c). At this stage, the solution of the continuation problem is said to be robust to new data and the continuation algorithm can perform a new prediction step. As the continuation algorithm progresses in parameter space, more data points are added to the GPR model. To keep computational costs low, an overall maximum number of data points n_{\max} in the model is maintained by removing data points that have less influence on the zero-problem.

During the experiment, the set of candidate data points is usually limited to 50 points. This is a much smaller set of points than the one considered for the colour maps in Figure 4. This explains why in Figure 4(b) there exist a small difference between the apparent location of the maximum of Eq. (10) and the location where the new data point has been collected (\times).

Other approaches to decide where to collect data points could have also been used. For instance, the effect of new data points on the solution of the zero-problem was investigated. However, this approach was found to give similar results to the method above while needing a solution to the nonlinear continuation problem for each candidate data point and hence being computationally much more expensive. Another approach was to select the data points for which the variance, $\text{var}[\Gamma_*]$, of the predicted distribution was the largest.

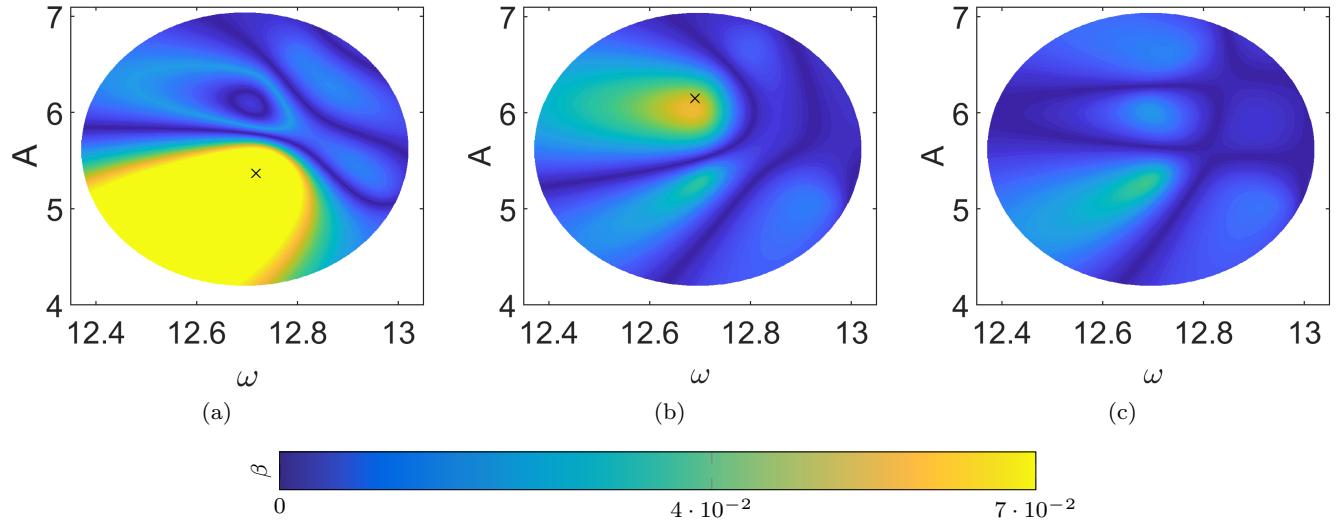


Fig. 4 Illustrates the new data selection process. The colourmap shows the sensitivity of the zero-problem β as estimated by Eq. (10). (a – c) Show the evolution of Eq. (10) when new data points (\times) are added to the GPR model. 4×10^{-2} is the user-defined threshold below which no additional data point is collected.

This approach was however discarded as new data points were being positioned at the periphery of the data set.

3 Description of the experimental set-up

The experimental set-up considered for the experimental demonstration of our algorithm is shown in Figure 5. The main structure is a steel ($\rho \approx 7850 \text{ Kg/m}^3$, $E \approx 210 \text{ GPa}$) cantilever beam of length 380 mm, width 31.9 mm and thickness 1.9 mm. The beam's free end is attached to two linear springs arranged as shown in Figure 5(b). This mechanism gives rise to geometric nonlinearity at large amplitudes. Previous work has shown that the stiffness properties of this mechanism can be approximated by a linear plus cubic term [30]. However, a mathematical model of the nonlinearity is unnecessary for CBC. As such, neither the identification of the nonlinear parameters nor the exploitation of the mathematical form of the nonlinearity were used. The length of the beam as well as the pre-tension in the springs were carefully adjusted such that the ratio between the natural frequencies of the first two bending modes is close to, but larger than 3. This leads to the presence of a 3:1 modal interaction between these modes.

The structure is excited approximately 40 mm away from the clamp using a GW-V4 Data Physics shaker powered by a PA30E amplifier. The force applied to the structure is measured using a PCB208C03 force transducer connected to a Kistler signal conditioner (type

5134). The vibrations of the beam are measured at the tip using an Omron ZX2-LD100 displacement laser sensor. The beam structure, the first laser sensor, the shaker and its power amplifier constitute the nonlinear experiment tested using CBC.

The algorithm used by the CBC method and presented in Section 2 is run on a laptop computer directly connected to the real-time controller (RTC) box via a USB cable. The RTC box consists of a Beagle-Bone Black on which the feedback controller used by CBC is implemented. Note that CBC algorithms do not run in real-time, only the feedback controller does. The BeagleBone Black is fitted with a custom data acquisition board (hardware schematics and associated software are open source and freely available [31]). All measurements are made at 1 kHz sampling with no filtering. Estimations of the Fourier coefficients of the response, input force, and control action are calculated in real time on the control board using recursive estimators [23]; however, this was for convenience rather than a necessity.

The z -domain transfer function of the controller used by the CBC technique is given by

$$\frac{U(z)}{E(z)} = \frac{0.0053}{z^3 - 2.4521z^2 + 1.9725z - 0.5155}, \quad (11)$$

and aims to reduce the error $E(z)$ between the beam tip response, y (laser 1), and a control reference signal, y^* , (see Section 4.1). The control law, which was found to stabilise the dynamics of the experiment through-

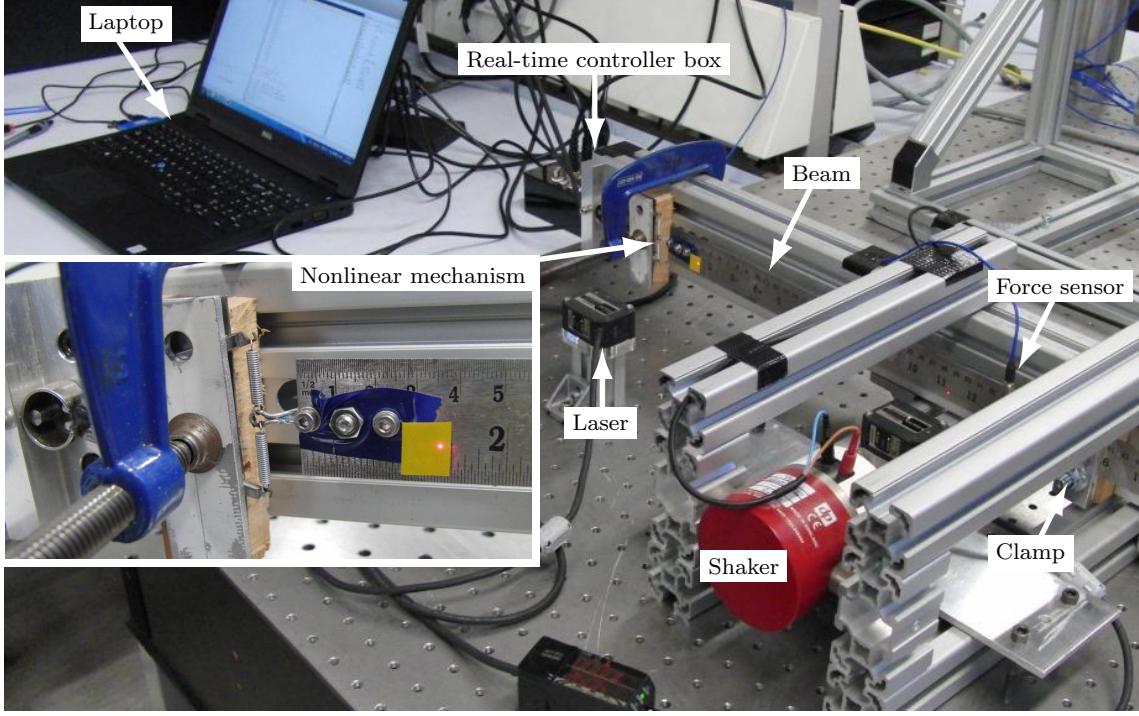


Fig. 5 Experimental set-up. The structure corresponds to a cantilever beam with a nonlinear mechanism attached at its free end. The structure is excited and controlled by means of an electro-dynamic shaker. A displacement laser sensor is used to measure the motion of the beam tip. CBC algorithms run on a laptop interconnected with the real-time controller box.

out the range of parameters considered in this study, was designed using pole-placement techniques and a linear model of the experiment. This model was obtained using low-amplitude broadband input-output data and captures the first two bending modes of the beam whose natural frequencies (damping ratios) were estimated at 11.49 Hz (0.026) and 36.45 Hz (0.022), respectively. Additional details on the derivation of the controller can be found in [21] where the nonlinear frequency response curves of the present experiment were also investigated using CBC. Note that errors in the model do not affect the results as long as the model is sufficiently accurate for designing a stabilising feedback controller.

4 Experimental results

The CBC technique used in this paper to collect experimental data points is briefly reviewed in Section 4.1 and exploited to map out the complete response surface of the system of interest. The new online regression-based algorithm is then demonstrated in Section 4.2 where it is used to track fold points while the experiment is running. To validate the new algorithm, obtained results are directly compared with the fold curves that can be extracted by post-processing the response surface ob-

tained using standard CBC approach as described in Section 4.1.

4.1 Extracting periodic responses

Considering periodic responses of the experiment, it is assumed that the control target, $y^*(t)$, and the control signal, $u(t)$, can be decomposed into a finite number m of Fourier modes as

$$y^*(t) = \frac{A_0^*}{2} + \sum_{j=1}^m A_j^* \cos(j\omega t) + B_j^* \sin(j\omega t), \quad (12)$$

$$u(t) = \underbrace{A_1^u \cos(\omega t) + B_1^u \sin(\omega t)}_{\text{fundamental}} \quad (13)$$

$$+ \underbrace{\frac{A_0^u}{2} + \sum_{j=2}^m A_j^u \cos(j\omega t) + B_j^u \sin(j\omega t)}_{\text{higher harmonics}}. \quad (14)$$

CBC seeks a reference signal $y^*(t)$ for which the control signal u is equal to zero for all times. When this is achieved, the control signal is non invasive and does not alter the position in parameter space of periodic responses compared to the underlying uncontrolled system of interest. To achieve this non-invasive control signal, the higher harmonics $(A_0^u, A_j^u, B_j^u)_{j=2}^m$ have to be eliminated by finding suitable reference signal coefficients $(A_0^*, A_j^*, B_j^*)_{j=2}^m$. This can be performed using

a derivative-free Picard-iteration algorithm. As for the fundamental harmonic component of the control signal, $A_1^u \cos(\omega t) + B_1^u \sin(\omega t)$, it can be viewed as the harmonic excitation applied to the experiment. If no other external excitation is applied to the experiment, the total harmonic excitation applied to the experiment is simply given by $\Gamma = \sqrt{A_1^{u^2} + B_1^{u^2}}$. The reader is referred to [19, 20] for a more detailed discussion on this approach that is often referred to as the simplified CBC method.

The higher-harmonic coefficients $(A_0^*, A_j^*, B_j^*)_{j=2}^m$ are determined to cancel out the higher-harmonic coefficients of the control signal. One of the fundamental target coefficients (here B_1^*) can also be set equal to zero in order to set a phase reference between the response and the excitation. As such, the only two adjustable inputs to the experiment are the frequency of excitation ω and the fundamental target coefficients A_1^* . The amplitude of the excitation Γ is viewed as a free parameter as it is not fully determined by the user but depends on the response $y(t)$ and target $y^*(t)$.

Considering a constant frequency, the response of the experiment to harmonic input can be investigated by increasing the free fundamental target coefficient A_1^* . The response curves collected in this way form curves that resemble a *S*, as shown in Figure 6. At each data point (●), full time series measurements containing 2000 samples, or between 22 and 29 oscillation cycles, were made and used to estimate response and force amplitudes. The number of Fourier modes considered in the calculation of amplitudes is equal to seven throughout the rest of the paper.

Response curves in Figure 6 were collected every 0.25 Hz between 10 Hz and 14 Hz with 0.2 mm steps in A_1^* . The *S*-curve obtained at 12.5 Hz is notably different from the other *S*-curves and includes an additional inflection point in its upper part. This curve lies in the region where a modal interaction occurs between the first two modes of the structure and marks a change in the form of the response surface of the system. This change in the response surface leads to the presence of an isola as further discussed in Section 4.2.

The *S*-curves collected can be post-processed to extract fold curves which will then be compared to the fold curves directly tracked during the experiment. The post-processing approach used to extract these curves is similar in principle to the one used in the online algorithm but with the clear distinction that it does not influence the data points collected. A GPR model in-

cluding all collected data points is created and numerical continuation is exploited to follow the solutions of Eq. (1). Figure 6 shows in gray the continuous surface constructed from the GPR model. The darker gray region indicates where the underlying uncontrolled experiment is unstable and thus unobservable without stabilising feedback control. The boundary of that region is the fold curve found using continuation. This curve is sensitive to the spacing between the different *S*-curves as well as the choice of hyper-parameters for the GPR. In particular, the fold curve presents artifact oscillation as discussed later. The online algorithm will overcome these issues.

4.2 Demonstration of the online regression-based algorithm

The online regression-based continuation algorithm presented in Section 2 is now demonstrated. Figure 7 illustrates the first four steps of the algorithm which is initialized in the neighbourhood of a fold found at high response amplitudes on one of the *S*-curves collected in Section 4.1 (○). Around that starting point, 25 new data points (●) regularly distributed in frequency and control target amplitude are collected. This first data set is then used to estimate the hyper-parameters of the GPR model using the marginalization approach presented in Section 2.3. Hyper-parameters were found to be $\theta = (\sigma_n^2, \sigma_f^2, l_\omega, l_A) \approx (0.02, 2.02, 0.30, 1.09)$. The noise variance, σ_n^2 , which is found to be relatively small compared to the other hyper-parameters, captures the effects of measurement noise. It also encompasses potential inaccuracies that could arise from the inaccurate cancellation of the control signal higher harmonics during data collection.

After estimation of the hyper-parameters, the first GPR model is created (■) and a first solution to Eq. (1) is found (×). The sequence of prediction and correction steps of the numerical continuation algorithm can then be started. Data points collected during the continuation process are shown in green (○) and appear to be regularly distributed around the fold of the response surface (see two-dimensional projections in Figure 7(b, c)). This comes from the regular pattern used for the generation of candidate data points and the sensitivity criterion (10) that is found to almost always have a structure similar to the one observed in Figure 4(a, b). Note that to collect data points at the particular response amplitudes specified by the data selection method in Section 2.5, a second regression mapping A to A_1^* is created during the continuation. This mapping

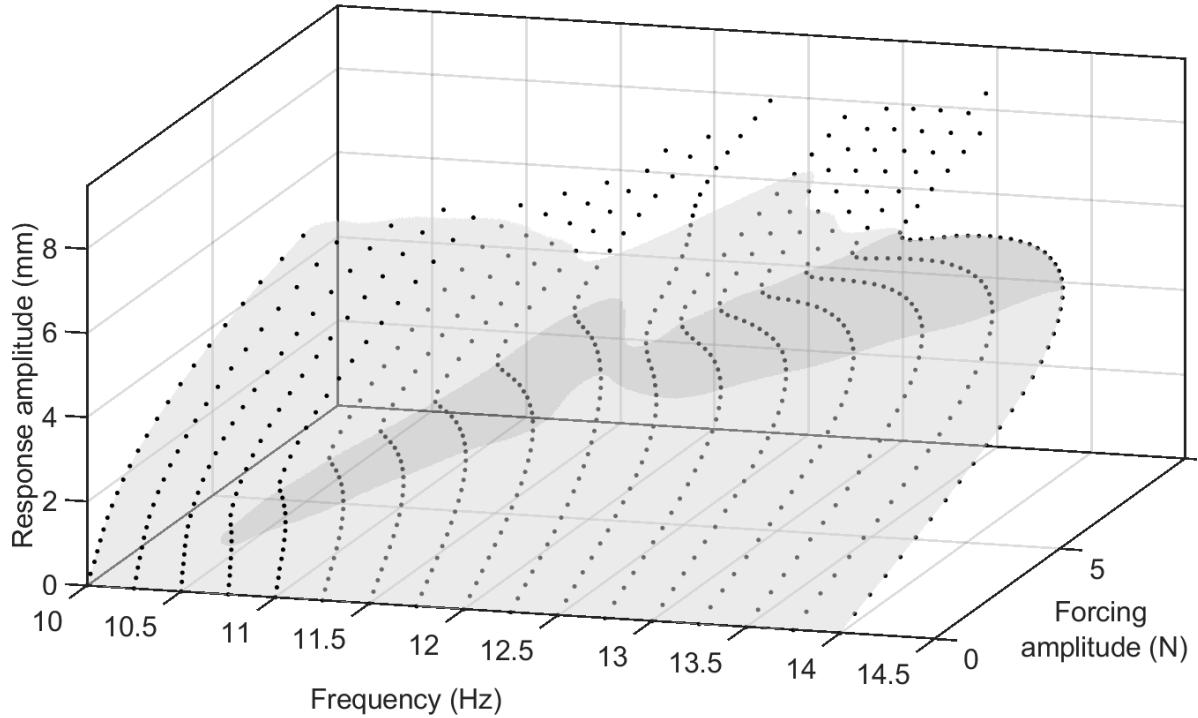


Fig. 6 *S*-curves collected using the simplified CBC algorithm presented in Section 4.1. Gray surface is the output of a GPR model including all recorded data points (●). The darker gray region is delimited by a fold curve. Data points in this region correspond to unstable response of the underlying uncontrolled experiment.

is observed to be almost linear throughout the range of parameters considered in this study (Figure 8).

In Figure 7, the red curve (—●—) represents the solutions found by the continuation algorithm and the force amplitude predicted at the solution by the GPR model. For each solution point, a data point is also collected (—●—) and added to the GPR model. At the beginning of the algorithm (around 14 Hz), the forcing amplitude estimated by the GPR model at the solution is slightly different from the force amplitude actually measured. However, the GPR model predictions improve as more data points are added to it and the difference between the two curves (—●—, —●—) becomes increasingly smaller, and negligible after 3 continuation steps.

Figure 9 illustrates the results obtained when applying the continuation algorithm throughout the parameter range of interest. All the data points collected during the continuation are shown in black (●), and the curve of fold points obtained by following the curve previously shown in Figure 7 is in red (—). Two-dimensional projections of the fold curve are shown in Figure 9(b, c). The curve has a complex geometry; especially in the region between 12 – 13 Hz where it is observed to increase in response amplitude at almost constant forcing

frequencies (Figure 9(b)). This feature corresponds to the ‘saturation’ of the resonance frequency commonly observed in nonlinear frequency response curves in the vicinity of modal interactions [32]. Closer inspection at this region reveals that the fold curve forms the shape of a swallowtail catastrophe (Figure 9(c)) [33]. Although this feature is small and appears to be more affected by measurement uncertainties than the rest of the curve, it was consistently observed in this region of the parameter space. After this interaction region, the algorithm was found to reach and successfully pass through a cusp (around 11 Hz).

The use of multivariate regression combined with a predictor-corrector continuation technique makes the online regression-based algorithm more general and robust than the bifurcation tracking algorithm proposed in [20]. This latter standard method could not have tracked a curve that exhibits the complex features of the fold curve in Figure 9. The present algorithm was however not able to systematically continue the fold curve across the modal interaction region. In particular, when the interaction region was approached from below, the fold curve was found to come very close to a previously measured section of the curve. Within uncertainty, this close proximity between the two portions of the curve gives rise to a branch point singularity which

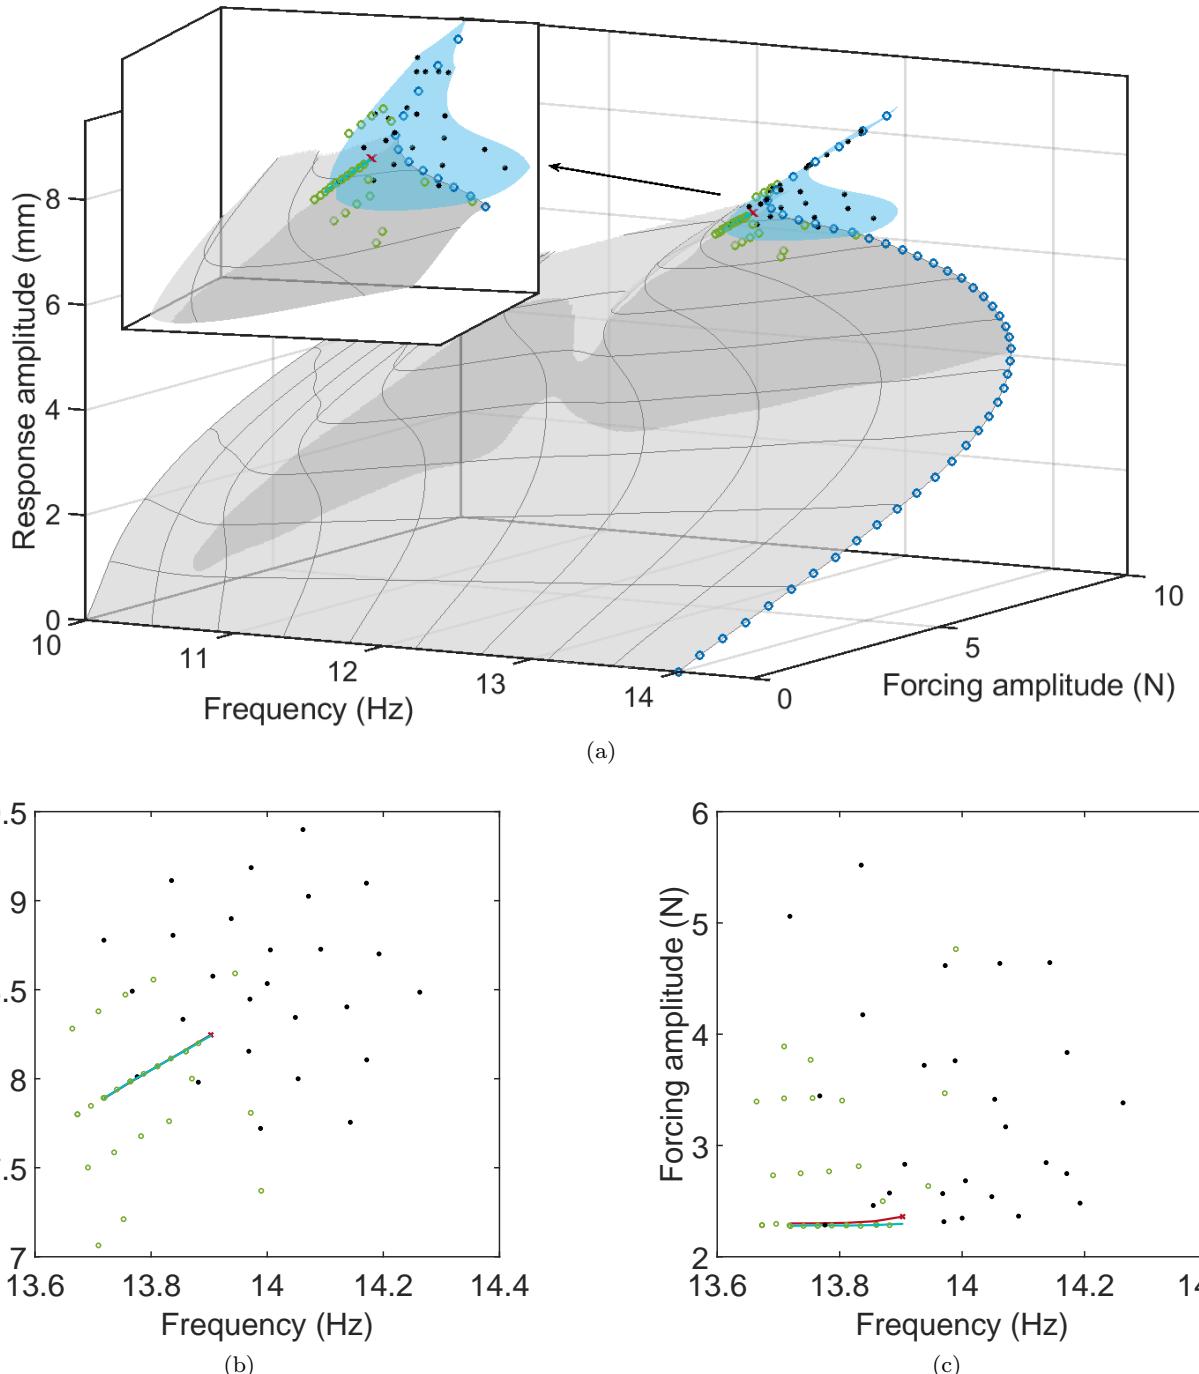


Fig. 7 First four steps of the regression-based continuation algorithm. (○) S-curve used to initialize the algorithm close to a fold. Data points (●) used to estimate hyper-parameters and find the first fold point. Data points (○) collected during the continuation process. (—●—) and (—●—) are the fold curves found using continuation and where the force amplitude is estimated by the GPR model and measured experimentally, respectively. (■) First local GPR model. To help visualization, the response surface obtained in Figure 6 is superimposed to the data obtained with the new algorithm. (b, c) Show two-dimensional projections of the data points and curves shown in (a).

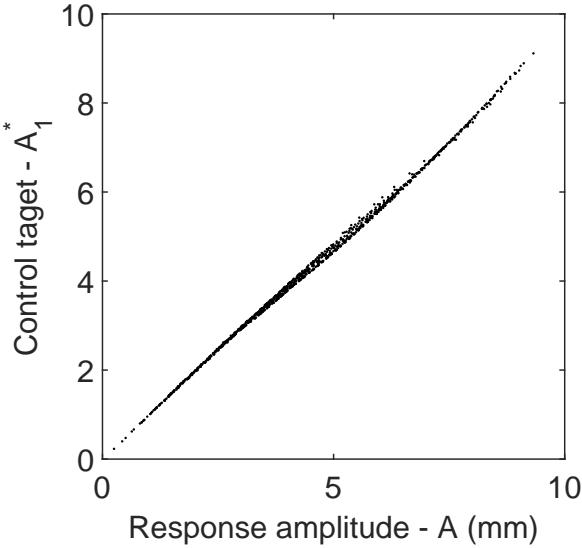


Fig. 8 Mapping between the response amplitude A (continuation variable) and the fundamental target coefficient A_1^* (experimental input). All the data points (\bullet) presented in this paper have been included.

cannot be addressed by our current continuation algorithm. As such, the online algorithm was often found to ‘jump’ to and continue the upper part of the fold curve. To capture the low-amplitude section of the fold curve located between 12.5 Hz and 14 Hz (— in Figure 9), the online algorithm was restarted around the second fold point observed at low amplitude on the S -curve in Figure 7.

To validate the continuation results, fold points found at a particular force amplitude are compared to the nonlinear frequency response (NLFR) curve obtained for the same excitation level. To generate the NLFR curve, the raw data points shown in Figures 6 and 9 that correspond approximately to the chosen force amplitude (within $\pm 5\%$) are represented on the same graph. The NLFR curves obtained for 2.0 N and 2.9 N are shown in Figures 10(a) and 10(b), respectively. Fold points are in red (\times). At 2N (Figure 10(a)), the resonance peak leans towards higher frequencies which highlights the hardening character of the nonlinearity in the system. The locations of the two fold points found at that level match the folds in the NLFR. Data points located between them are observed to be unstable periodic responses of the underlying uncontrolled experiment. As the excitation is increased to 2.9 N (Figure 10(b)), the resonance peak is qualitatively unchanged but an isolated branch of high-amplitude responses disconnected from the main resonance peak has appeared [30, 21]. The presence of this isola could have been inferred from Figure 9(c) as a third fold point ap-

pears when the force amplitude becomes larger than approximately 2.1 N.

To further validate the fold curves obtained using the online regression-based algorithm, results are compared to the fold curve obtained by post-processing the data points of the full response surface. In particular, a GPR model including all the data points of Figure 6 is created, and numerical continuation is applied to find and track the solutions of (1). Two sets of hyper-parameters are considered to post-process the data. The first one corresponds to the set of hyper-parameters used by the online algorithm. The second one is obtained by maximising the marginal likelihood for all the data points of the response surface. Hyper-parameters found in this latter case are $(\sigma_n^2, \sigma_f^2, l_\omega, l_A) \approx (0.01, 2.66, 0.28, 0.73)$. Although different from the hyper-parameters found at the initialization of the online continuation algorithm, they are found to have comparable magnitudes.

The resulting fold curves are compared in Figure 11. The two curves (—, —) obtained by post-processing the response surface are almost identical, which shows that the difference between the two sets of hyper-parameters does not play a significant role here. These curves are found to oscillate around the fold curve measured online (— \bullet —). These oscillations are artifact created by the regression method due to the lack of data points (in particular at resonance between successive S -curves). These oscillations are not present in the curve obtained directly during the experiment due to the tailored collection of suitable data points during the continuation. The blue curve is also found to reproduce the swallowtail catastrophe observed on the data collected online. The post-processing applied to the data points of the response surface could have also been applied to the data points collected by the online regression-based continuation algorithm. The fold curve obtained in this way (not shown for conciseness) is a smooth and continuous curve that is no longer affected by the ‘jump’ observed during the experiment and that also reproduces the swallowtail catastrophe. The curve is also free from regression artifacts as collected data points were suitably chosen.

To understand the sensitivity of the results to the selection of data points collected, 300 fold curves were numerically computed for different data sets. Each set is based on all the data points collected using the online algorithm but with 10% of the points, selected randomly, removed. The hyper-parameters of the GPR model were marginalized for each data set. Fold curves

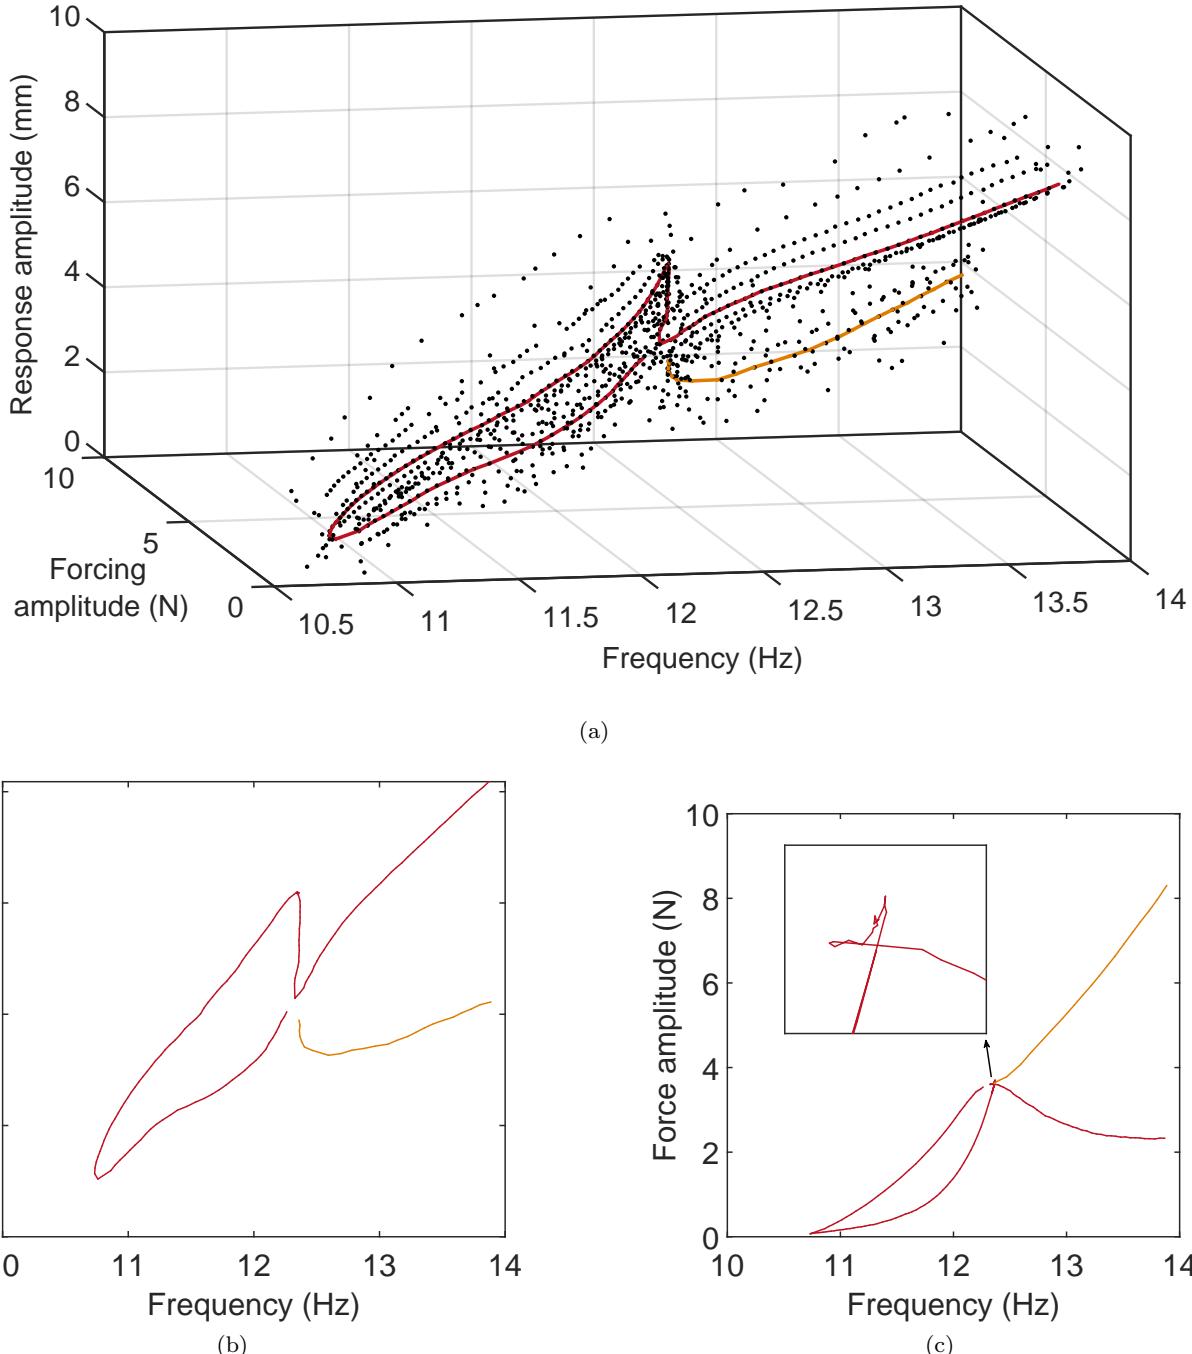


Fig. 9 (—, —) Fold curves captured experimentally using the regression-based continuation algorithm. (●) Data points automatically collected during continuation using the method proposed in Section 2.5. (b, c) Two-dimensional projections of the fold curves shown in (a). The close-up in (c) shows the presence of a swallowtail catastrophe.

(—) are shown in Figure 12. Regions of larger result variability, i.e. where the overall bifurcation curve appears thicker, are located on the lower part of the curve (Figure 12(a)). This is expected as this region of the response surface is significantly less curved than the resonance region and hence more affected by measurement

uncertainties. Large variability is also noticeable at the cusp as well as at the modal interaction region. In the latter, the distinction between the upper and lower part of the fold curve can disappear as some fold curves are found to cross and have a branch point. However, other continuation runs are also found to lead to disjoint fold

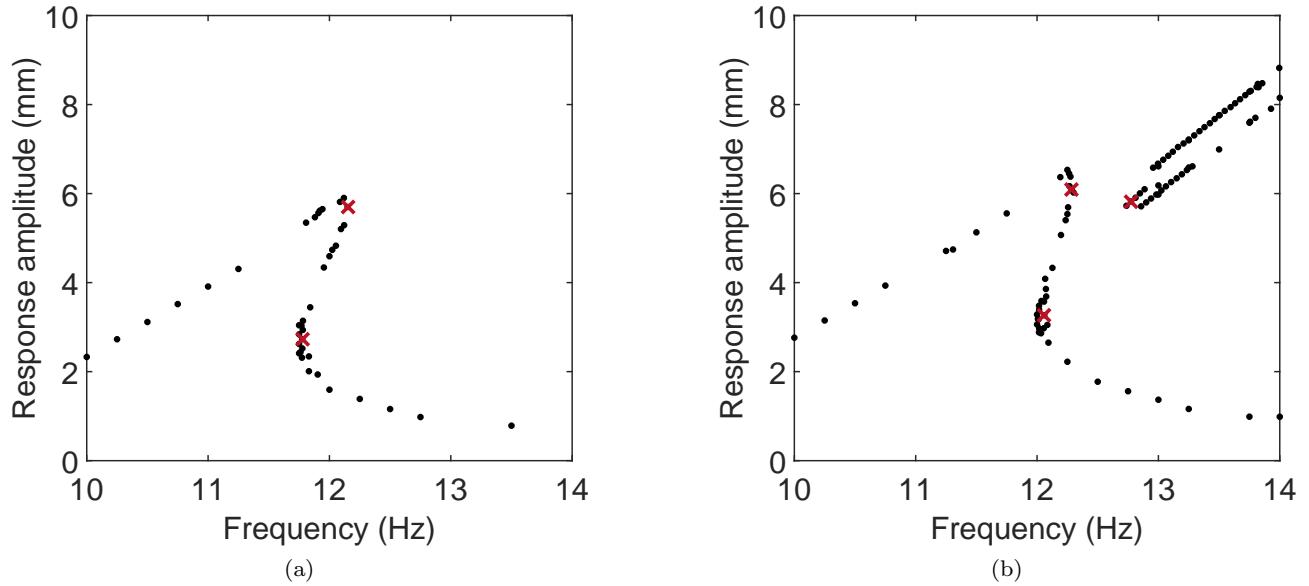


Fig. 10 Nonlinear frequency response of the system obtained by selecting data points in Figures 6 and 9 that correspond approximately (within $\pm 5\%$) to (a) 2.0 N and (b) 2.9 N. Fold points (\times) obtained with the regression-based continuation algorithm of Section 2.

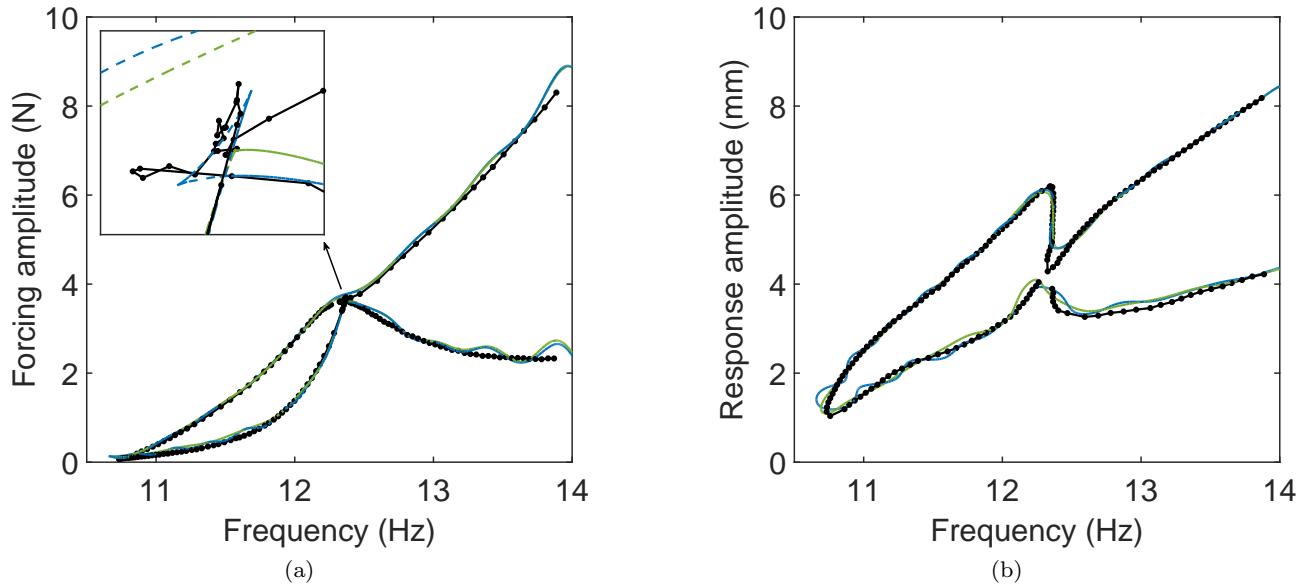


Fig. 11 Comparison between the fold curve ($-\bullet-$) obtained during the experiment and the fold curves ($-\text{---}$) obtained by post-processing the data points of the response surface. Hyper-parameters used to obtain ($-\text{---}$) are identical to those used in the experiment, whereas the hyper-parameters used for ($-\text{---}$) result from the maximisation of the marginal likelihood for all the data points in the response surface.

curves — one before and one after the modal interaction. This important variability in the results highlights the “flat” nature of the response surface in that region and hence the difficulty to infer the presence of folds (inset in Figure 12(a)). In general, the presence of uncertainty may make it difficult to draw strong conclusions on the exact nature of the dynamics in such a

region. Figure 12(b) shows the sensitivity of the results with respect to the bifurcation parameters forcing frequency and amplitude. As observed experimentally, the swallowtail catastrophe appears to be a robust feature of the system’s dynamics as it is present for almost all data sets.

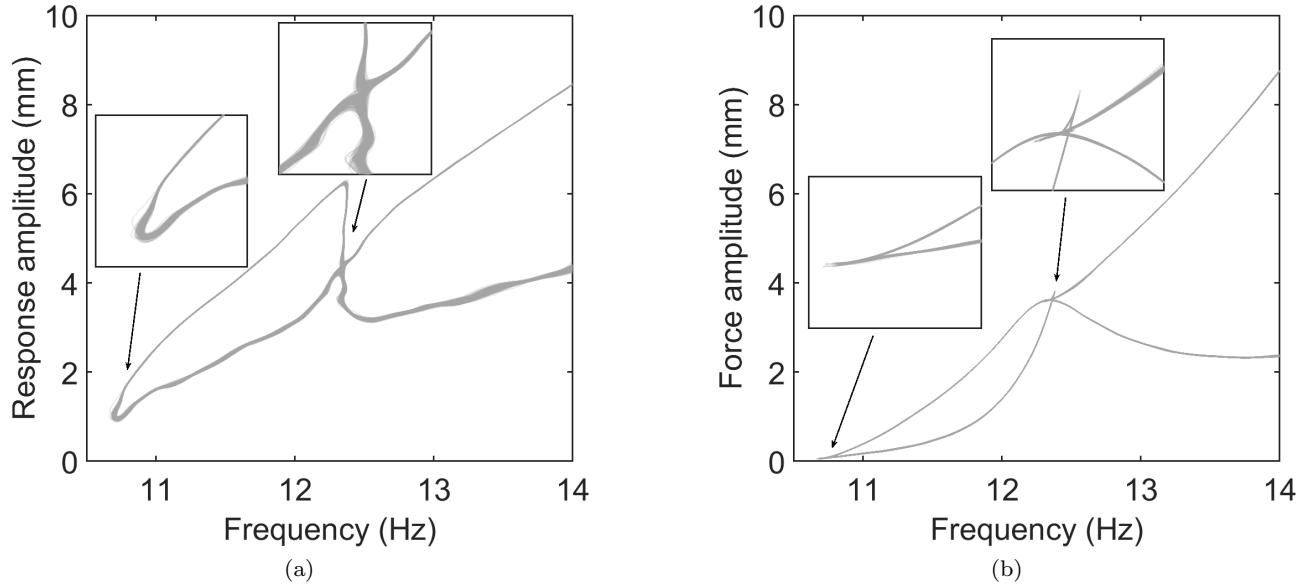


Fig. 12 300 fold curves (—) obtained by post-processing some of the data points measured during the online continuation. For each curve, 10% of the data points are randomly selected and removed from the data set. Each curve is also computed with hyper-parameters values obtained via the maximization of the marginal likelihood for the particular data set considered.

5 Conclusions

Control-based continuation is a general and systematic method to probe the dynamics of nonlinear experiments. In this paper, CBC is combined with a novel continuation algorithm that is robust to experimental noise and enables the tracking of geometric features of the response surface. The online regression-based CBC uses Gaussian process regression to create local models of the response surface. These models are smooth (noise-free) which enables the use of standard numerical continuation algorithms. An important aspect of the proposed method is to interact with the experiment to select the data points to collect such that they are the most informative about the dynamic feature of interest. This approach is similar in principle to an experimental design problem in statistics and could a priori be applied to other problems such as nonlinear parameter estimation and model selection.

The online regression-based algorithm was experimentally demonstrated on a nonlinear structure with harmonically-coupled modes by tracking fold points in the structure response to harmonic excitation. The algorithm was able to address the complex geometry of the fold curve arising from the presence of a modal interaction. However, the close proximity between two portions of the fold curve lead, within uncertainties in the experiment, to a branch point singularity which could not be addressed by the current continuation al-

gorithm.

The online regression-based algorithm presented in this paper is very general and could be exploited to capture other types of dynamic features such as backbone curves (nonlinear normal modes) and nonlinear frequency response curves. The regression algorithm used here relies however on a unique parameterisation of the solution curve in terms of the continuation variables. If this assumption is not satisfied, further improvements of the method to consider local coordinates for the regression, or a larger set of continuation variables might be necessary.

Data Statement

Experimental data collected in this study are available at [DOI to be inserted at proofing].

Acknowledgements L.R. has received funding from the Royal Academy of Engineering, fellowship RF1516/15/11. D.A.W.B. is funded by the EPSRC grant EP/K032738/1 and S.A.N. by the EPSRC fellowship EP/K005375/1. We gratefully acknowledge the financial support of the Royal Academy of Engineering and the EPSRC.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

- Y. A. Kuznetsov. *Elements of Applied Bifurcation Theory*, volume 112 of *Applied Mathematical Sciences*. Springer New York, 2004.
- R. Seydel. *Practical Bifurcation and Stability Analysis*, volume 5 of *Interdisciplinary Applied Mathematics*. Springer New York, 3rd. edition, 2010.
- S. Godwin, D. Ward, E. Pedone, M. Homer, A.G. Fletcher, and L. Marucci. An extended model for culture-dependent heterogenous gene expression and proliferation dynamics in mouse embryonic stem cells. *npj Systems Biology and Applications*, 3(1):19, 2017.
- B. Krauskopf, H. M. Osinga, E. J. Doedel, M. E. Henderson, J. Guckenheimer, A. Vladimirsky, M. Dellnitz, and O. Junge. A survey of methods for computing (un)stable manifolds of vector fields. *International Journal of Bifurcation and Chaos*, 15(3):763–791, 2005.
- M. Peeters, R. Viguié, G. Sérandon, G. Kerschen, and J. C. Golinval. Nonlinear normal modes, part II: Toward a practical computation using numerical continuation techniques. *Mechanical Systems and Signal Processing*, 23(1):195–216, 2009.
- S. Huntley, D. Jones, and A. Gaitonde. Bifurcation tracking for high reynolds number flow around an airfoil. *International Journal of Bifurcation and Chaos*, 27(04):1750061, 2017.
- E. Ott, C. Grebogi, and J.A. Yorke. Controlling chaos. *Physical Review Letters*, 64:1196–1199, 1990.
- S. Misra, H. Dankowicz, and M.R. Paul. Event-driven feedback tracking and control of tapping-mode atomic force microscopy. *Proceedings of the Royal Society A*, 464:2113–2133, 2008.
- K. Pyragas. Delayed feedback control of chaos. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1846):2309–2334, 2006.
- R.E. Best. *Phase-Locked Loops — Design, Simulation and Applications*. Mc Graw Hill, 6th Edition.
- S. Mojrzisch, J. Wallaschek, and J. Bremer. An experimental method for the phase controlled frequency response measurement of nonlinear vibration systems. In *Proc. Appl. Math. Mech*, volume 12, pages 253–254, 2012.
- S. Peter and R. I. Leine. Excitation power quantities in phase resonance testing of nonlinear systems with phase-locked-loop excitation. *Mechanical Systems and Signal Processing*, 96:139–158, 2017.
- V. Denis, M. Jossic, C. Giraud-Audine, B. Chomette, A. Renault, and O. Thomas. Identification of nonlinear modes using phase-locked-loop experimental continuation and normal form. *Mechanical Systems and Signal Processing*, 106:430–452, 2018.
- J. Sieber and B. Krauskopf. Control based bifurcation analysis for experiments. *Nonlinear Dynamics*, 51(3):365–377, 2008.
- J. Sieber, B. Krauskopf, D. Wagg, A. Gonzalez-Buelga, and S. Neild. Control-based continuation of unstable periodic orbits. *Journal of Computational and Nonlinear Dynamics*, 6(1):011005, 2010.
- E. Bureau, F. Schilder, I. Ferreira Santos, J. J. Thomsen, and J. Starke. Experimental bifurcation analysis of an impact oscillator - tuning a non-invasive control scheme. *Journal of Sound and Vibration*, 332(22):5883–5897, 2013.
- E. Bureau, F. Schilder, M. Elmegrd, I. F. Santos, J. J. Thomsen, and J. Starke. Experimental bifurcation analysis of an impact oscillator - determining stability. *Journal of Sound and Vibration*, 333(21):5464 – 5474, 2014.
- M. Elmegård, B. Krauskopf, H. M. Osinga, J. Starke, and J. J. Thomsen. Bifurcation analysis of a smoothed model of a forced impacting beam and comparison with an experiment. *Nonlinear Dynamics*, 77(3):951–966, 2014.
- D. A. W. Barton and J. Sieber. Systematic experimental exploration of bifurcations with noninvasive control. *Physical Review E*, 87(5):052916, 2013.
- L. Renson, D. A. W. Barton, and S. A. Neild. Experimental tracking of limit-point bifurcations and backbone curves using control-based continuation. *International Journal of Bifurcation and Chaos*, 27(01):1730002, 2017.
- L. Renson, A.D. Shaw, D.A.W. Barton, and S.A. Neild. Application of control-based continuation to a nonlinear structure with harmonically coupled modes. *Mechanical Systems and Signal Processing*, 120:449 – 464, 2019.
- D. A. W. Barton and S. G. Burrow. Numerical continuation in a physical experiment: Investigation of a nonlinear energy harvester. *Journal of Computational and Nonlinear Dynamics*, 6(1):011010, 2010.
- L. Renson, A. Gonzalez-Buelga, D. A. W. Barton, and S. A. Neild. Robust identification of backbone curves using control-based continuation. *Journal of Sound and Vibration*, 367:145–158, 2016.
- L. Renson, D.A.W. Barton, and S.A. Neild. Experimental analysis of a softening-hardening nonlinear oscillator using control-based continuation. In *Proceedings of the International Modal Analysis Conference (IMAC)*, Orlando, USA, 2016.
- F. Schilder, E. Bureau, I. F. Santos, J. J. Thomsen, and J. Starke. Experimental bifurcation analysis - continuation for noise-contaminated zero problems. *Journal of Sound and Vibration*, 358:251–266, 2015.
- G. Kerschen, K. Worden, A. F. Vakakis, and J. C. Golinval. Past, present and future of nonlinear system identification in structural dynamics. *Mechanical Systems and Signal Processing*, 20(3):505–592, 2006.
- J.P. Noël and G. Kerschen. Nonlinear system identification in structural dynamics: 10 more years of progress. *Mechanical Systems and Signal Processing*, 83:2–35, 2017.
- C. E. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, Massachusetts, 2006.
- A. McHutchon and C. E. Rasmussen. Gaussian process training with input noise. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, NIPS’11, pages 1341–1349, 2011.
- A. D. Shaw, T. L. Hill, S. A. Neild, and M. I. Friswell. Periodic responses of a structure with 3:1 internal resonance. *Mechanical Systems and Signal Processing*, 81:19–34, 2016.
- D. A. W. Barton. Real-time control hardware/software based on the beaglebone black, 2015.
- C. Chen, D.H. Zanette, D.A. Czaplewski, S. Shaw, and D. Lpez. Direct observation of coherent energy transfer in nonlinear micro-mechanical oscillators. *Nature Communications*, 8(15523), 2017.
- T. Poston and I. Stewart. *Catastrophe Theory and Its Applications*. Dover, 1978.