Mechanical Systems and Signal Processing 70-71 (2016) 527–541

Contents lists available at ScienceDirect

Mechanical Systems and Signal Processing

journal homepage: www.elsevier.com/locate/ymssp

Chatter detection in turning using persistent homology

Firas A. Khasawneh [a,*], Elizabeth Munch [b]

[a] Department of Mechanical Engineering, SUNY Polytechnic Institute, 100 Seymour Rd, Utica, NY 13502, United States
[b] Department of Mathematics & Statistics, University at Albany – SUNY, 1400 Washington Ave, Albany, NY 12222, United States

ARTICLE INFO

Article history:
Received 10 November 2014
Received in revised form
22 September 2015
Accepted 27 September 2015
Available online 23 October 2015

Keywords:
Machining dynamics
Nonlinear delay differential equation
Stochastic delay equation
Topological data analysis
Persistent homology

ABSTRACT

This paper describes a new approach for ascertaining the stability of stochastic dynamical systems in their parameter space by examining their time series using topological data analysis (TDA). We illustrate the approach using a nonlinear delayed model that describes the tool oscillations due to self-excited vibrations in turning. Each time series is generated using the Euler-Maruyama method and a corresponding point cloud is obtained using the Takens embedding. The point cloud can then be analyzed using a tool from TDA known as persistent homology. The results of this study show that the described approach can be used for analyzing datasets of delay dynamical systems generated both from numerical simulation and experimental data. The contributions of this paper include presenting for the first time a topological approach for investigating the stability of a class of nonlinear stochastic delay equations, and introducing a new application of TDA to machining processes.

1. Introduction

Deterministic models for chatter in machining dynamics have been the subject of extensive research in recent years [1]. However, machining processes are inherently stochastic with many noise sources due, for example, to the variation of the material parameters [2,3], external noise sources [4], and variations of the delay [5]. Thus, the scarcity of analysis tools for stochastic delay systems remains a fundamental impediment to the continuous progress in the metal removal industry [6,7].

Stochastic equations are infinite dimensional and therefore analysis tools from deterministic models are not readily applicable to them. The analysis is more challenging if the dynamics involve delays so that the system model is a stochastic delay differential equation SDDE). Since machine tool chatter is typically described by these difficult systems, the number of studies on stochastic machining dynamics remains small, particularly in comparison to its deterministic counterpart. We note here that in addition to machining dynamics [3,4], SDDEs arise in many applications such as chemical kinetics [8] and genetic networks [9]. Therefore, developing or extending analytical and numerical tools for their analysis continues to be an active and important area of research.

For a limited number of SDDEs, stochastic calculus can be used to study the stability of the first and second moments [10]. If the delay is small, then the SDDE can be approximated using a stochastic differential equation without the delay term [11]. An extension of the semi-discretization method for studying the moment stability of linear SDDEs with delays appearing in the drift term only was described in [12]. Another method to investigate the stability of this class of equations

* Corresponding author.
E-mail address: firas.khasawneh@sunyit.edu (F.A. Khasawneh).

http://dx.doi.org/10.1016/j.ymssp.2015.09.046
0888-3270/

uses the Lyapunov approach [13]. However, for the general case of SDDEs numerical simulation remains the most viable method of analysis.

Euler–Maryuama and Milstein simulation methods, originally used to simulate stochastic differential equations, were extended to SDDEs in [14–17] and [18], respectively. Numerical simulation provides a tool for generating path-wise solutions that can be easier to investigate than the original SDDE. For example, instead of directly studying the mean square (or more generally, the pth mean) stability of the SDDE which might be difficult or impossible, the paths generated by numerical simulation can be used [19–21]. The result of the numerical simulation is a time series or dataset that contains information about the system dynamics. Developing data analysis tools for these datasets has two benefits: (1) it provides a benchmark for testing new methods for the analysis of SDDE, and (2) the same tools can be used to analyze data from real-world applications, e.g. machining dynamics.

Some of the data analysis methods for non-delayed stochastic equations include principal component analysis [22,23], multi-dimensional scaling [24], local linear embedding [25], Laplacian eigenmaps [26], Hessian eigenmaps [27], local tangent space alignment [28], and diffusion maps [29–31]. The first step in many of these methods is to obtain a lower-dimensional representation of the underlying high-dimensional manifold. The hope is that the simplified representation captures the main features of the underlying dynamics. One of the key assumptions in many of the prominent data analysis methods for dynamical systems, e.g. diffusion maps, is that the underlying dynamics is Markovian. This precludes them from being used to study SDDEs which are non-Markovian.

Frequency domain methods, typically based on Fourier transform and the power spectrum, have also been used for data analysis especially when the signal is periodic. However, when the signal is non-periodic, the leakage effects inherent in Discrete Fourier Transform (DFT) will result in errors unless the input is periodic and the length of the time series is equal to the input's period [32]. Further, in the context of machining we are interested in changes in the time series associated with the change of the cutting process from stable to unstable; however, frequency domain analysis does not account for changes in the signal with time. Other issues with the frequency domain analysis include the limitations in analyzing trending time series and data with quasi-periodic motion [33]—a common route to chaos which is often associated with unstable cuts [34]. In order to address these limitations, some success has been reported using time-frequency methods [35] and wavelets [36]. However, the topic of analyzing stochastic machining models remains an active area of research.

In this paper we explore data analysis tools for studying the stability of a stochastic turning model using topological data analysis. These tools are applicable to datasets arising from both experiments as well as simulations of dynamical systems. Specifically, we will use persistent homology [37–40] to automatically detect when changes in the system behavior indicative of chatter occur near the stability boundary of the linearized, noise-free model. In contrast to other data analysis tools, persistent homology does not attempt to obtain a lower dimensional representation of a data set, but rather a low-dimensional descriptor which is easy to understand and which can be used to find and measure properties of interest.

Persistent homology has found success in applications to many diverse fields such as neuroscience [41,42], genetics [43,44], epidemiology [45], tracking [46,47], international relations [48], map reconstruction [49], sensor networks [50,51], and image analysis [52]. Most recently, a great deal of work has looked at using persistence for signal analysis. The idea behind this method is to use the Takens embedding [53] to turn a signal into a point cloud in high dimensional space and analyze the resulting point cloud using persistence. In [54], it was shown that this method provides a framework to parameterize dynamics; the procedure is also amenable to combination with machine learning for automation [55]. Variations of this idea along with broader applications of topology have been used to study signals from human speech [56], wheezing in breathing signals [57], gene expression data [58–61], computer architecture [62,63], and character animation [64]. See also Chap. 6 of [65] or [66] for the formulation of the procedure using cohomology.

The standard Takens theorem assumes that the system is deterministic and autonomous [53]. However, the theorem has been extended to deterministic non-autonomous equations and to stochastic systems [67–69], as well as to infinite-dimensional dynamical systems [70]. Since the focus of this paper is on describing a topological approach for studying nonlinear stochastic machining models, we utilize the standard Takens theorem with the understanding that embedding timesets with large stochastic terms is likely to be unsuccessful. The extension of Takens theorem to nonlinear delay models with large stochastic terms is outside the scope of this paper and it is a topic of active research.

We demonstrate the main concepts using a second order nonlinear stochastic delay equation with multiplicative noise that models a single degree of freedom turning process. To keep the analysis clear, we only introduce one random component using a stochastic cutting force coefficient where the source of stochasticity can be the variation in the temperature, shear angle, or workpiece material properties. In particular, we study the model SDDE in the space of the non-dimensional spindle speed and depth of cut, and we use Euler–Maruyama method to simulate the SDDE for different combinations of these two process parameters. We show that as the value of the delay is varied as a result of varying the spindle speed, persistent homology can be used to detect the change of the response from a steady state equilibrium to a periodic orbit indicating the loss of stability through a Hopf bifurcation. We also demonstrate the effect of the noise intensity on the stability diagrams.

2. Background

In this section, we will cover the background necessary to describe the method of the experiment. We give brief introductions to persistent homology (Section 2.1); to Takens embedding theorem (Section 2.2); and to the relationship between periodicity and the maximum persistence (Section 2.3).

2.1. Persistent homology

We begin with an informal introduction to the subject of persistent homology, and direct the reader to [71,72] for a full introduction to classical homology and to [73,40,39] for a more in-depth introduction to persistence. Suppose we are given a point cloud drawn from a manifold and want to understand something about the underlying structure. To do this, we consider expanding a collection of discs of the same radius centered at each point. We can then study the structure of the union of these discs for a changing radius. In the example of Fig. 1, we start with a point cloud sampled from an annulus and see that at a very small radius (around ($r=0.3$)), the collection of purple disks consists of a set of disconnected components. At a slightly larger radius ($r=0.07$), these discs start to intersect, possibly forming small circular structures which fill in at a still slightly larger radius. What is very interesting is that at a relatively small radius (approximately $r=0.9$), we form a circular structure representing the full annulus which takes a much longer time to fill in than the small circular structures seen previously. The goal of persistent homology is to quantify this intuition in a rigorous manor and use it to "measure" the size of the annulus.

Next, we represent the union of discs in a combinatorial structure which is amenable to computation. This structure is called a simplicial complex and can be thought of as a generalization of a graph. Whereas a graph has only nodes and edges, a simplicial complex has nodes, edges, triangles, tetrahedra, and higher dimensional analogues, called simplices. A simplex σ in the simplicial complex is represented by a subset of the nodes. In particular, given a set of points $\chi \subset \mathbb{R}^n$ and $a \geq 0$, we approximate the structure of the union of discs by a particular simplicial complex called the Rips complex, $\mathcal{R}_a$. This is a simplicial complex which consists of a vertex v_i associated to each point $x_i \in \chi$, and has an edge (v_i, v_j) if the corresponding points are within distance a of each other, $\|x_i - x_j\| \leq a$. Then higher dimensional simplices are added whenever possible; that is, the simplex σ is in $\mathcal{R}_a$ iff $\|x_i - x_j\| \leq a$ for all $v_i, v_j \in \sigma$.

The Rips complex for a particular a can be thought of as an approximate representation of the union of disks for radius $a/2$ since an edge is added between two points at distance a exactly when two disks of radius $a/2$ touch. In the example of Fig. 1, we draw edges to represent the Rips complex using $a=2r$ for the labeled radius r. In addition, for visualization we

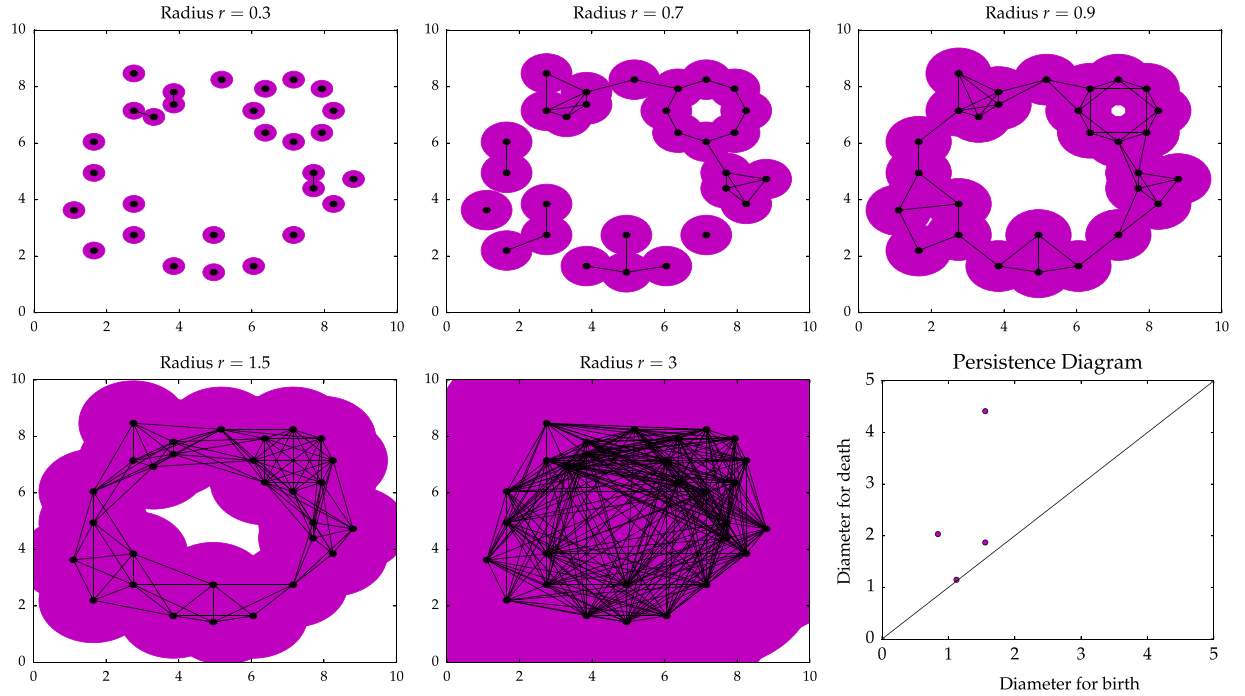

Fig. 1. A point cloud shown with collections of purple disks for several different radii are shown in the first five figures; the persistence diagram for this point cloud is given in the sixth figure. The Rips complexes, $\mathbb{R}_{2r}$, are drawn on top of the point clouds by drawing the edges, and then assuming that every triangle which appears is filled in. The circular structure of the point cloud is reflected in the single point far from the diagonal in the persistence diagram. The distance from this point to the diagonal is called the maximum persistence and can be used to quantify the circular structure.

only draw edges, then assume that we have included the triangles and tetrahedra whenever they appear. Squares as in the bottom left of the radius 0.9 figure will not be filled in.

Now that we have a way to represent the union of discs, we need a method for quantifying the changing structure seen. This tool comes by way of homology, a construction from classical algebraic topology [71,72]. The idea is to represent a particular type of structure of the space in a vector space, called the p-dimensional homology group and denoted $H_p(\mathbb{X})$ when computed for a space $\mathbb{X}$. Note that we can always assume the homology groups are vector spaces since we work with $\mathbb{Z}_2$ coefficients. In this work, we focus on the 1-dimensional homology group, $H_1(\mathbb{X})$, which measures how many loops are seen in the space. In particular, for a circle, the first homology group is a vector space of rank one, thus telling us that the circle has one loop. For a Fig. 8 structure, the rank of the first homology group is two, representing the two loops in the space. When we obtain these spaces from some sort of embedding in Euclidean space, such as when we take the circle or Fig. 8 and draw them in the plane, homology is rotation and translation invariant; so the homology of 8 and ∞ (as topological spaces) are the same. This paper focuses on studying the first homology group for the Rips complex at various diameters a_i. For example, the rank of $H_1(\mathcal{R}_{0.6})$, given the point cloud in Fig. 1, is zero because there are no holes present in the picture for radius $r=0.3=a/2$. The ranks of the first homology groups for $r=0.7, 0.9, 1.5$, and 3 in the example are 2, 3, 1, and 0 respectively.

There is, however, quite a bit more information that we can take from these vector spaces. In particular, not only do we have an understanding of the homology at a fixed diameter a, but we can relate these vector spaces to each other as the diameter parameter changes. Specifically, if $r\leq s$, we have an inclusion $f:\mathcal{R}_r\longrightarrow\mathcal{R}_s$, and this induces a linear map on the vector spaces, $f_*^{r,s}:H_p(\mathcal{R}_r)\rightarrow H_p(\mathcal{R}_s)$. Given any sequence of increasing $0\leq a_1\leq a_2\leq\cdots\leq a_n$, we have a *filtration*, $\mathcal{R}_0\subseteq\mathcal{R}_{a_1}\subseteq\cdots\subseteq\mathcal{R}_{a_n}$. Often, we choose the diameters $\{a_i\}$ to be the collection of pairwise distances between points $\{\|v_i-v_j\|\}_{i,j}$, since these are exactly the times that the structure of the union of discs changes. Then we can collect these inclusions into a sequence of inclusions of simplicial complexes, called a *filtration*, written as

$$\mathcal{R}_{a_1}\longrightarrow\mathcal{R}_{a_2}\longrightarrow\cdots\longrightarrow\mathcal{R}_{a_n}.$$

Each of these inclusions gives a map on homology which can also be collected together in sequence

$$H_p(\mathcal{R}_0)\longrightarrow H_p(\mathcal{R}_{a_1})\longrightarrow H_p(\mathcal{R}_{a_2})\longrightarrow\cdots\longrightarrow H_p(\mathcal{R}_{a_n}).$$

This construction is called a *persistence module*. Through a theorem of Carlsson and de Silva [74], the information in a persistence module can be uniquely represented by a *persistence diagram*. This is a collection of pairs (b,d), where each pair corresponds to the diameter b when a feature appears, and the diameter d when that same feature disappears; the full persistence diagram is often denoted D. In the last image in Fig. 1, we have drawn the 1-dimensional persistence diagram for the given point cloud. Note that because there is a large circular structure that appears (is born) at a small diameter and gets filled in (dies) at a much larger diameter in the union of disks, there is a point far from the diagonal in the persistence diagram. Small loops that appear and disappear quickly in the set of expanding discs appear as points close to the diagonal. In this way, we have a visual representation of the difference between classes which live a long time (and could be considered important), and the points representing short lived classes (and are often assumed to be noise). For this reason, we usually draw the diagonal $\{(x,x)|x\in\mathbb{R}\}$ into the persistence diagram along with the off-diagonal points to remind the observer that points always appear above the diagonal, and that proximity to the diagonal implies that an off-diagonal point is more likely to be noise.

2.2. The Takens embedding

The Takens embedding is a standard tool for time series analysis [33]. Given a time series $X(t)$, which in practice is a set of samples $s_n=X(t_n)$, fix a time lag $\eta>0$ and choose a dimension $m\in\mathbb{Z}_{>0}$ in which to embed the data. Then the Takens embedding is a lift of the time series to the map to $\mathbb{R}_m$

$$\psi_\eta^m: t\longmapsto(X(t),X(t+\eta),\ldots,X(t+(m-1)\eta)).$$

Through an important theorem of Takens [53], we are justified in using the term "embedding" since, under the correct parameter choices, this mapping preserves the structure of the underlying manifold and the dynamics of the system.

Of course, this embedding is sensitive to the choices of the parameters m and η. A choice of m which is too low means that the structure of the embedding will not accurately reflect the dynamics; a high choice of m increases computation time due to the curse of dimensionality. Additionally in the discrete setting, a large choice of m will also amplify noise since a single value in the time series affects m points in the embedding. With respect to the lag parameter, a low choice of η means that the data will be overly correlated and the data will cluster around the diagonal in $\mathbb{R}_m$, yielding no information. On the other hand, a high choice of η gives an embedding which will be too spread out to say anything useful.

Luckily, methods exist which can help to intelligently choose the embedding parameters. We used false nearest neighbors [33] to determine embedding dimension m. This procedure works by finding "false" neighbors of a particular point which are only close because of projections caused by an embedding dimension which is too low, and looking for a high enough dimension to get rid of the majority of these. In order to choose the time lag parameter η, we used the autocorrelation function $c(\eta)$ [33]. The function $c(\eta)$ gives information about how the points are distributed at the choice of lag η. If they cluster around the diagonals $X(t_i)=X(t_i-\eta)$ or $X(t_i)=-X(t_i-\eta)$, the autocorrelation function will be positive or

negative, respectively. If, however, the points are evenly spread out, $c(\eta) \approx 0$. Thus, we choose η to be the first zero of the autocorrelation function.

2.3. Persistence and amplitude

Persistent homology, in some sense, measures the size of holes in manifolds. This study is concerned with quantifying periodicity in a time series, which translates to a loop in the Takens embedding. Hence, we can compute the 1-dimensional persistence of the generated point cloud in order to measure the size of this structure. An excellent study of the relationship between a periodic signal and the maximum persistence can be found in [58].

First, consider the three signals in the first row of Fig. 2. These are example signals from the experiment discussed in Section 5. The first consists of only noise, the second has periodic behavior with a small amplitude, and the last has periodic behavior with a large amplitude. The Takens embedding for each is drawn in the second row; notice that because we decided on the delay parameter η for each time series individually using the autocorrelation function, η is different for each diagram. The periodic signals have embeddings with a circular structure, while the noisy signal has a point cloud which is a dense blob. In order to quantify this observation, we can look at the persistence diagram for each, shown in the last row.

The noticeable difference between these diagrams is the appearance of an off-diagonal point which is far from the diagonal. So, for this project, we will compare the persistence diagrams using the maximum persistence, defined as

$$\text{MaxPers}(D) = \sup_{(b,d) \in D} |d-b|$$

for any diagram D. In this example, the first persistence diagram has $\text{MaxPers}(D) \approx 1.5$, the second has $\text{MaxPers}(D) \approx 0.25$, while the maximum persistence for the last diagram is nearly 0. So, the more periodic the signal, the more circular the point cloud, the larger the maximum persistence.

This procedure is most useful because of its resilience to noise. A simple corollary to the stability theorem of Cohen-Steiner et al. [75] shows that if we have two point clouds which are close (here we mean in the Hausdorff distance), then the difference in the computed maximum persistence will be even smaller. It should also be noted that the larger the amplitude, the larger the hole in the circular structure, and thus the larger the maximum persistence. For these reasons, we use maximum persistence to quantify periodicity.

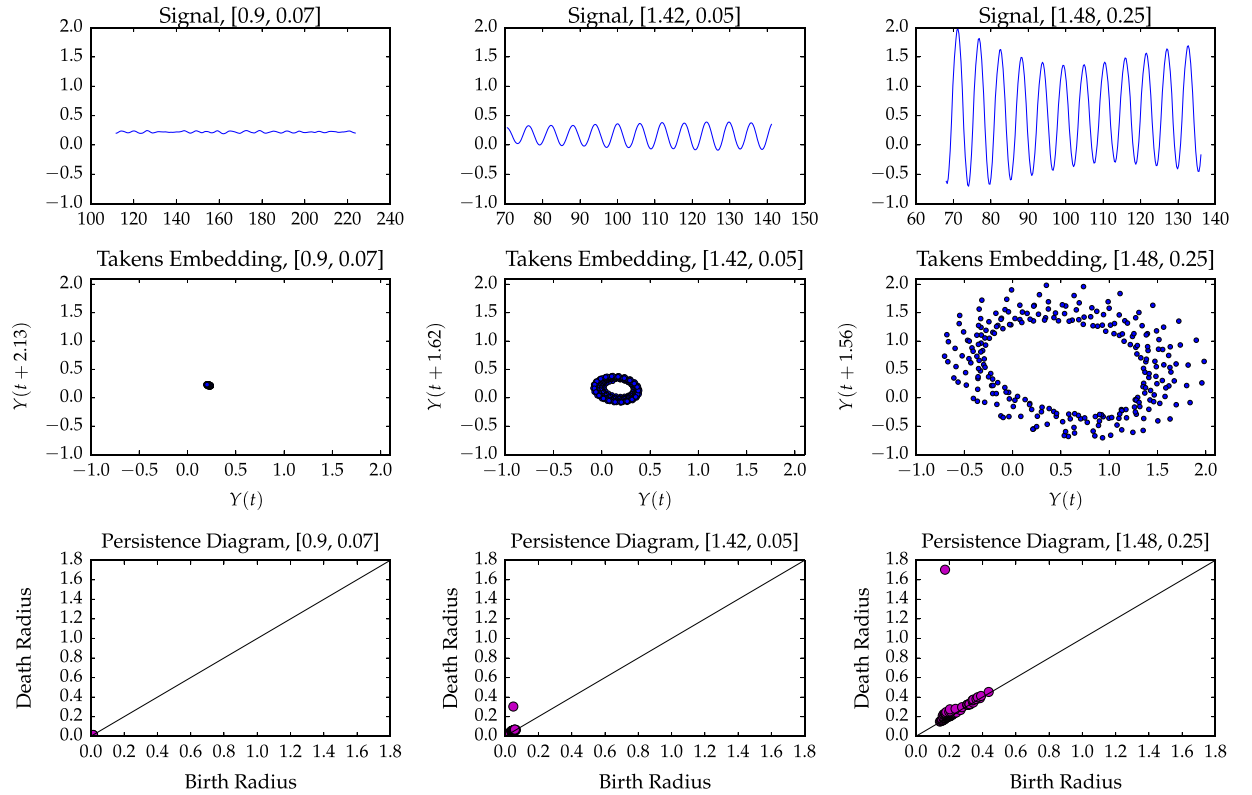

Fig. 2. This example shows the relationship between periodicity and maximum persistence. The first row contains three example signals, averaged from 20 realizations with $\delta = .01$ as discussed in Section 5 from parameter locations shown in Fig. 4 by the circle, triangle, and X, respectively. The second row shows the Takens embedding of each, and the last row shows the respective 1-dimensional persistence diagrams with the diagonal marked. Note that the periodic signals have a prominent off-diagonal point in the persistence diagram, while the non-periodic signal has no such point.

3. Mathematical model

Fig. 3 shows the turning process that will be investigated in this study. In this figure, the tool is modeled as a linear oscillator with a single degree of freedom in the y direction while the workpiece is assumed to be rigid. The cutting force F acts at the tool tip and the resulting equation of motion reads

$$\ddot{y}+2\zeta\omega_n\dot{y}+\omega_n^2 y=\frac{F}{m} \tag{3.1}$$

where ζ is the damping ratio, ω_n is the natural angular frequency of the tool, and m is the modal mass. We use the power model for the cutting force [76,77] which is assumed to depend on the uncut chip area A shown in Fig. 3 according to

$$F=Kwh^{\alpha} \tag{3.2}$$

where K is the mechanistic cutting coefficient with units of $\mathrm{N/m^2}$ which relates the cutting force to the area of the uncut chip $A=w\times h$, w is the depth of cut shown in Fig. 3, h is the chip thickness, and the exponent is typically chosen as $\alpha=0.75$.

Since the tool is a dynamic system, it will oscillate under the action of the cutting force and leave behind a wavy surface on the workpiece representing the history of the tool oscillations during one workpiece revolution. The tool will then advance into the workpiece according to a set rate per spindle revolution h_0 with units of m/rev, called the nominal feed rate, where it will encounter the undulated surface that it left behind during the previous cut. Therefore, the uncut chip thickness will vary dynamically because it depends on the current oscillations of the tool, represented by $y(t)$ in Fig. 3, and the oscillations left behind one spindle revolution earlier, represented by $y(t-\tau)$ in the same figure. Here, $\tau=2\pi/\Omega$ is the time it takes to complete one spindle revolution corresponding to the spindle speed Ω rad/s. This coupling between the cutting force and the current as well as the delayed oscillations of the tool can lead to self-regenerative vibrations which can become unstable, thus giving rise to chatter.

If the tool oscillations are too large, then the tool will lose contact with the workpiece and the instantaneous chip thickness will become zero since there is no material being cut. So, the instantaneous chip thickness can be written as

$$h(t)=\begin{cases} h_0+y(t-\tau)-y(t) & \text{if } y(t)-y(t-\tau)\le h_0 \\ 0 & \text{otherwise.}\end{cases}$$

Note that the instantaneous chip thickness includes the contribution of the nominal feed rate h_0 as well as the current and the delayed tool oscillations given by $y(t)$ and $y(t-\tau)$, respectively.

To reduce the number of parameters in Eq. (3.1), we use the rescaling described in [78]: let $y(t)=h_0\tilde{y}(t)$ and rescale time such that $\tilde{t}=\omega_n t$ and $\tilde{\tau}=\omega_n\tau$. After dropping the tildes for the sake of notation, the resulting equation reads

$$\ddot{y}+2\zeta\dot{y}+y=\frac{Kw(2\pi R)^{\alpha-1}}{m\omega_n^2}\rho^{\alpha-1}(1+y(t-\tau)-y(t))^{\alpha}=b\rho^{\alpha-1}(1+y(t-\tau)-y(t))^{\alpha} \tag{3.3}$$

where R is the radius of the workpiece, $\rho=h_0/(2\pi R)$, and b is the dimensionless depth of cut. Typical values for ρ in conventional turning are $\rho<0.01$. The corresponding condition for the tool to be in contact with the workpiece is

$$y(t)-y(t-\tau)\le 1. \tag{3.4}$$

From Eq. (3.3) it can be seen that the constant steady state solution is $b\rho^{\alpha-1}$. Eq. (3.3) and condition (3.4) are used in the simulation to generate time series for the deterministic system.

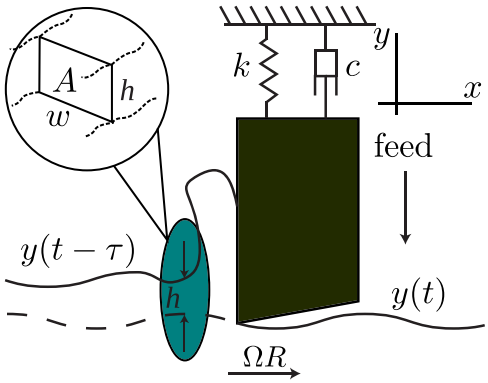

Fig. 3. The turning process under study in this paper. The tool is compliant in y while the workpiece is assumed rigid. The parameter Ω represents the spindle speed, R is the radius of the workpiece, h is the instantaneous chip thickness, w is the chip width, while $y(t)$ and $y(t-\tau)$ represent the tool oscillations during the current and the previous periods, respectively.

If we allow the dimensionless cutting coefficient to become a stochastic variable due to variations in the workpiece material, shear angle, or temperature effects, then we can write the stochastic non-dimensional cutting coefficient $\hat{b}$ as

$$\hat{b} = \overline{b} + \delta \frac{dB}{dt} \tag{3.5}$$

where $\overline{b}$ is the average or nominal value of the stochastic non-dimensional cutting coefficient $\hat{b}$, δ is the diffusion coefficient, and B is standard Brownian motion. The coefficient δ represents a measure of the stochasiticity in the system. In this context, it specifies the inhomogeneity in the material which, for instance, can be the result of strain hardening or manufacturing imperfections. As $\delta \to 0$, $\hat{b} \to \overline{b}$ and the behavior of the stochastic system becomes more and more deterministic. To elaborate, as δ approaches 0, the cutting coefficient approaches the nominal value $\overline{b}$ thus modeling the workpiece as a homogeneous material. Using the above definition, the stochastic delay differential equation describing the tool oscillations is

$$d\dot{Y} = \left(-2\zeta\dot{Y} - Y + \overline{b}\rho^{\alpha-1}(1 + Y(t-\tau) - Y(t))^{\alpha}\right)dt + \delta\left(\rho^{\alpha-1}(1 + Y(t-\tau) - Y(t))^{\alpha}\right)dB \tag{3.6}$$

where Eq. (3.6) is interpreted in the Itô sense [79].

Let $h = (1 + Y(t-\tau) - Y(t))$. We can write Eq. (3.6) as a pair of first order equations for the position and velocity of the tool according to

$$\begin{aligned} dY &= d\dot{Y}dt, \\ d\dot{Y} &= \left(-2\zeta\dot{Y} - Y + \overline{b}\rho^{\alpha-1}h^{\alpha}\right)dt + \delta\left(\rho^{\alpha-1}h^{\alpha}\right)dB. \end{aligned} \tag{3.7}$$

Eq. (3.7) is a nonlinear stochastic delay differential equation with multiplicative noise. In this paper, we use it in an Euler–Maruyama simulation to generate the noisy time series in Section 5.

4. Numerical simulation

Two different simulations were used: an explicit Runge–Kutta triple via Matlab's DDE23 command for the deterministic case in Eq. (3.3), and an Euler–Maruyama simulation for the stochastic system of Eq. (3.7). For the stochastic case, multiple realizations of the solution path were generated for use in the analysis. The specifics of the experiments are discussed in more detail in Section 5.

The system parameters that were used in the simulations are $\zeta = 0.03$, $\rho = 0.01$, and $\alpha = 0.75$. The simulation was performed over a 100×100 grid in the $(\Omega/\omega_n, b)$ space. Since the oscillations of the tool depend on the past oscillations, the simulation must include a record of the surface left behind by the tool during previous spindle revolutions. This record is called the history function The simulations were initialized by setting this history function equal to the steady state solution $b\rho^{\alpha-1}$ for $t \in [-\tau, 0]$ and a small perturbation (we chose 0.01) was introduced at $t = 0$ as shown in Fig. 5b. The information obtained from the simulation was then used to update the profile of the workpiece in subsequent revolutions as shown in Fig. 5c.

4.1. Stability of the deterministic linearized model

To provide a basis of comparison for the results obtained using the described approach and to elucidate the effect of noise on the system stability, we also studied the linearized stability analysis of the noise-free model (3.3). Eq. (3.3) describes the oscillations of the tool assuming a deterministic model and is used for obtaining the stability of the linearized deterministic system. Specifically, linearizing Eq. (3.3) around the steady state solution yields the linear delay differential equation

$$\ddot{\xi} + 2\zeta\dot{\xi} + y = \alpha b\rho^{\alpha-1}(\xi(t-\tau) - \xi(t)). \tag{4.1}$$

Eq. (4.1) is a linear autonomous delay differential equation whose stability can be studied using several methods in the literature, e.g. zero-order approximation [81], semi-discretization [82], Chebyshev collocation [83], or spectral element method [80]. In this study, we use the spectral element method with one temporal element, a polynomial of degree 12, and a 100×100 grid in the $(\Omega/\omega_n, b)$ plane, where Ω is the spindle speed given in rad/second, ω_n is the natural frequency of the tool, and b is the non-dimensional depth of cut introduced in Eq. (3.3). We suppress the details of the spectral element method and instead refer the reader to [80] for a more thorough description.

The resulting stability diagram is shown in Fig. 4. Shaded regions are stable while unshaded regions are unstable. The results are in agreement with the semi-discretization solution in [78]. Note that this line is drawn on the results shown in Fig. 8, however it is provided only as a means for comparison.

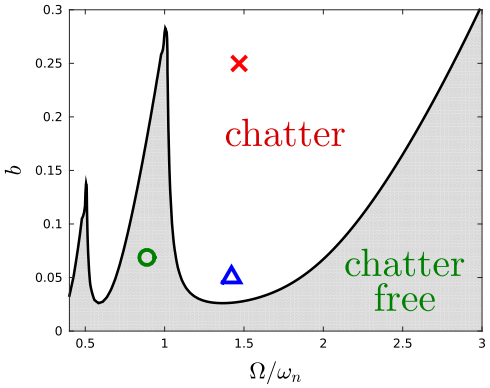

Fig. 4. The stability boundary for the linearized non-dimensional deterministic model (4.1) of the turning process calculated using the spectral element method [80] is shown as the black line in the figure. The parameter values in the shaded region correspond to stable cutting, while the unshaded region corresponds to unstable cutting. The circle, triangle, and X correspond to parameter choices for the signal examples given in Figs. 2 and 6.

4.2. Deterministic case simulation

Eq. (3.3) was simulated using Matlab's DDE23 function. The condition of Eq. (3.4) was monitored and the simulation was terminated as soon as the tool lost contact with the workpiece. The system was simulated over $t \in [0, 200]$, however the time series corresponding to unstable cuts were shorter because the exit of the tool from the cut caused early termination of the simulation.

4.3. Stochastic case simulation

Eq. (3.6) was simulated using the Euler–Maruyama method, described in [14], to generate the datasets. The Brownian path was created using Matlab and the approach described in [84]. For a fixed parameter pair $(\Omega/\omega_n, b)$, the time series was generated using Matlab's command rng('shuffle') which seeds the random number generator based on the current time. The equation was simulated over $[0, T]$ where $T = M \cdot \tau$ and we chose $M = 2^5$. Therefore, as Ω/ω_n varied, the end time of the simulation varied according to

$$T = M \cdot \tau = M\frac{2\pi}{\Omega/\omega_n}. \tag{4.2}$$

The number of points for generating the Brownian path is $N = M \cdot 2^{14}$ and the time steps for the Brownian path and for the Euler–Maruyama simulation are dt and Δt, respectively. In this study, we set $\Delta t = dt = M\tau/N$. This choice of the simulation parameters enabled mapping the term $Y(t-\tau)$ exactly onto both a point on the Brownian path and a point that was simulated previously, thus eliminating the need for estimating intermediate values. The time for each spindle revolution, which is equal to the time delay, is an integer multiple of the simulation time step according to

$$\tau = \frac{N}{M}\Delta t. \tag{4.3}$$

This means that the profile left behind in any workpiece revolution is discretized using N/M points.

We considered three values for the noise intensity in Eq. (3.7): $\delta = 0.01$, 0.03, and 0.05. Condition (3.4) was monitored during the stochastic simulation; however, in contrast to the deterministic case, the simulation was not terminated when the tool left the workpiece. Instead, the forcing term was turned off and the tool was subject to free oscillations. This means that it was possible for the tool to go in and out of the cut multiple times during the duration of the simulation. Fig. 5 shows four snapshots of the tool related to the loss of contact. Snapshots 1 and 2 show the tool right before and right after it lost contact with the workpiece. Between snapshots 2 and 3 the tool experiences free damped oscillations outside of the cut until it returns to the cut in snapshot 3. Consequently, during the next revolution the tool encounters the surface that was not cut due to the loss of contact as shown by snapshot 4. The code kept track of all the loss of contact events and updated the N/M surface profile representation points left behind accordingly. Examples of the resulting series can be found in Fig. 6.

5. Experiments

We used persistent homology to analyze the time series generated by the numerical simulation. For every experiment, each time series $X_{\Omega/\omega_n,b}(t)$ was given as a series of samples $s_n = X_{\Omega/\omega_n,b}(t_n)$ for $t_n \in [0, T]$, where T depends on the value of Ω/ω_n as described in Eq. (4.2). Only the second half of the time series was retained in order to consider the portion of the time series after transience. Because the simulated time series was quite dense (see Eq. (4.3)), the shortened time series was

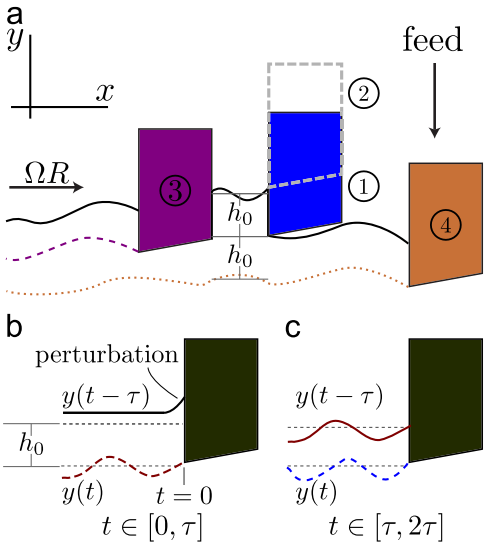

Fig. 5. Snapshots of the turning process. Each snapshot represents the tool at an instant in time during its constrained motion along the y-axis. (a) The surface profile is affected by tool oscillations along y. The tool is shown during four instants in time. At position 1, the tool is cutting the workpiece but due to the cutting forces it disengages the workpiece at time instant 2. It then oscillates freely (not shown) until time instant 3 where it regains contact with the workpiece until the end of the spindle revolution. Position 4 shows the time instant at the beginning of the next spindle revolution. The tool encounters the surface left behind during the previous revolution including any surface left behind due to loss of contact, e.g., between the time instants 2 and 3. (b) The history function during the first revolution which describes the workpiece profile for $t \le 0$ is chosen to be the equilibrium solution with a small perturbation at $t = 0$. (c) For subsequent revolutions (the second workpiece revolution is shown here), the surface profile is updated according to the tool oscillations in the previous revolution. As the workpiece spins, the surface profile left behind from previous spindle revolutions becomes the history function for the current revolution.

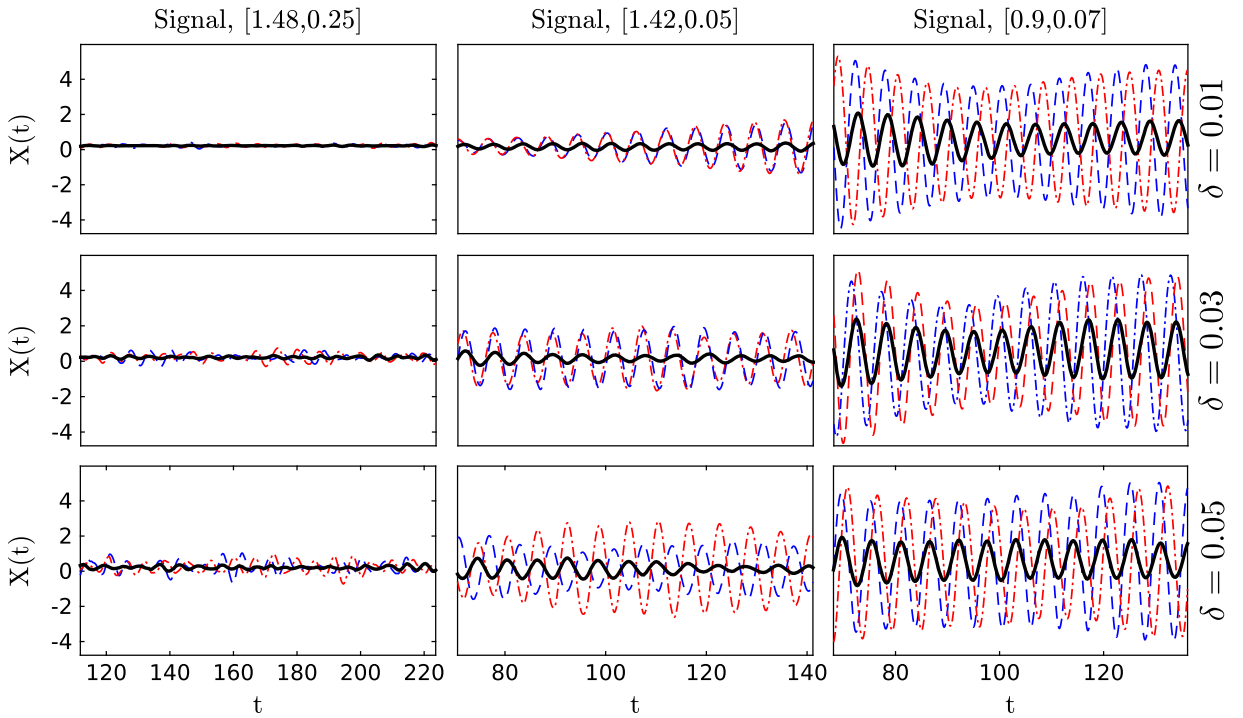

Fig. 6. Sample time series that show the (non-dimensional) oscillations of the cutting tool for three different combinations of the parameters Ω/ω_n and b written as a pair of values $(\Omega/\omega_n, b)$ in each column's heading. These three locations correspond to the points marked by a circle, triangle and an X in Fig. 4. Each column corresponds to one of the (Ω/ω_n) pairs while each row corresponds to a different value of the noise level δ. Within each figure we only show two sample realizations (dashed, colored) as well as the pointwise average of 10 realizations (solid, black). To highlight the differences between the time series, the scales of the y-axes for all the graphs are kept constant. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

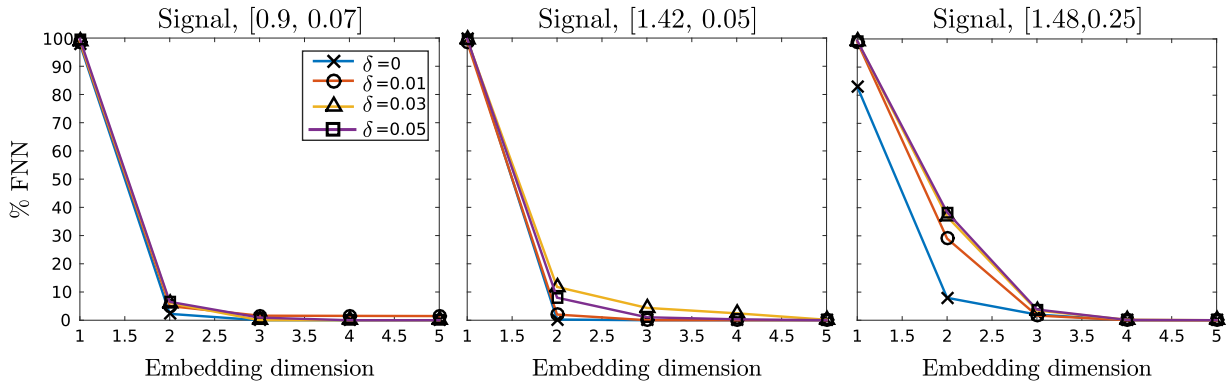

Fig. 7. The percentage of false nearest neighbors versus the embedding dimension calculated using the algorithm in [85]. Each graph corresponds to a pair of values $(\Omega/\omega_n, b)$ in the parameter space of the turning process as indicated in the graphs' headings. These three locations correspond to the points marked by a circle, triangle and an X in Fig. 4. Within each figure the percentage of the false nearest neighbors is plotted for increasing values of the noise level δ. For $\delta > 0$, performing the calculations on all the realizations produced very similar trends; therefore, we only show the results for one realization per value of δ here.

subsampled in order to speed the persistence computation. This resulted in 264 remaining data points evenly distributed in $[T/2, T]$ for all the time series both with and without noise.

The resulting time series was then embedded as a point cloud in $\mathbb{R}_3$. The embedding dimension $m = 3$ was determined using the standard false nearest neighbor algorithm [85]. Specifically, Fig. 7 shows that an embedding dimension of 3 is sufficient to reconstruct the phase space as evidenced by the drop of the corresponding percentage of the false nearest neighbors to less than 4.5% for all the considered cases. In the false nearest algorithm, the value of the threshold for designating a neighbor as a false neighbor was chosen to be $R_{\text{tol}} = 10$. We used the standard deviation of the time series as an estimate of the attractor size, and for the second criteria of the embedding algorithm for finding if the nearest neighbor is not necessarily close to a data point we picked the threshold $A_{\text{tol}} = 2$.

Each point cloud was reconstructed from the signal using a time lag η that was determined using the first zero of the autocorrelation function for that particular time series. Consequently, the point cloud consisted of points

$$\psi_\eta^3(t_n) = (X_{\Omega/\omega_n,b}(t_n), X_{\Omega/\omega_n,b}(t_n+\eta), X_{\Omega/\omega_n,b}(t_n+2\eta))$$

for $t_n \in [T/2, T-2\eta]$. Persistence was computed for this point cloud using the M12 package [86], and the maximum persistence of the 1-dimensional persistence diagram was saved.

The first experiment was designed to test the noise-free case. One time series was computed for each parameter pair $(\Omega/\omega_n, b)$ and maximum persistence of the Takens embedding was computed. The results of this experiment are shown in Fig. 8 (top right).

To study the effect of stochasticity on the results, the maximum persistence was computed for two collections of time series. First, for each parameter pair $(\Omega/\omega_n, b)$ and each noise value $\delta = 0.01, 0.03$, and 0.05, we generated ten realizations of the time series and computed the maximum persistence for each. The second row of Fig. 8 shows the average maximum persistence for these realizations for different choices of δ. Second, we generated twenty realizations of the time series for each $(\Omega/\omega_n, b)$ pair and computed the point-wise average $\overline{X}_{\Omega/\omega_n,b}(t)$ of these realizations. That is, if $X^i_{\Omega/\omega_n,b}$ represents the ith time series, then the resulting time series was

$$\overline{X}_{\Omega/\omega_n,b}(t) = \frac{1}{20}\sum_{i=1}^{20} X^i_{\Omega/\omega_n,b}(t). \tag{5.1}$$

Persistence was then computed for $\overline{X}_{\Omega/\omega_n,b}$ and the results are shown for the three noise values in the third row of Fig. 8. For easier comparison, the stability boundary for the deterministic case computed using the spectral element method is drawn using a white boundary line on all diagrams.

6. Results and discussion

The results of the described approach are shown in Fig. 8. It is important to notice that each row in Fig. 8 uses a different color scale. Since the simulation of the noise-free case is terminated as soon as the tool leaves the cut, the corresponding time series have a much smaller amplitude in comparison to the stochastic cases. In contrast, for the stochastic cases the simulation continues after the tool leaves the cut but the cutting forces are turned off and the tool's governing equation becomes that of an unforced harmonic oscillator. We believe the reason for the difference between the function values for the second (average of maximum persistence) and third (maximum persistence of the average) rows of Fig. 8 is caused by periodic behavior in the averaged time series which has a smaller amplitude because of the averaging process on

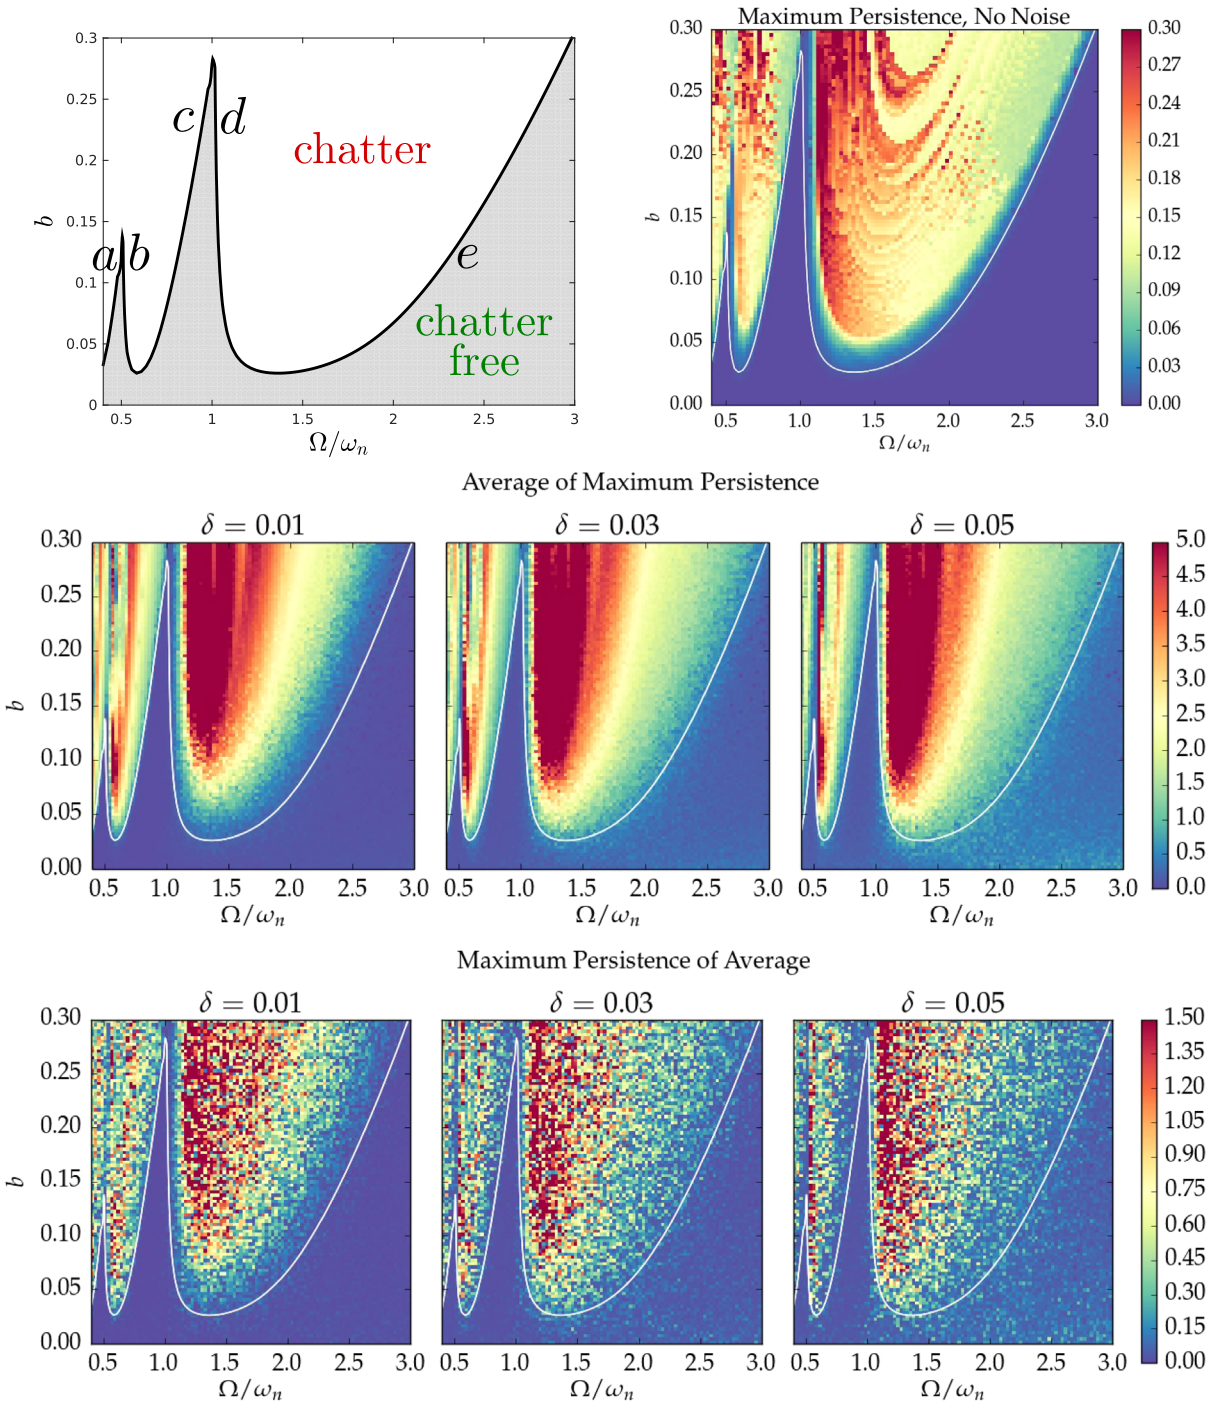

Fig. 8. Stability diagrams obtained using the described topological data analysis approach. The top left graph is a reproduction of Fig. 4 with some regions near the boundary labeled to facilitate the analysis in Sections 6 and 7. The top right graph is the stability diagram for the noise-free case. The second row shows the stability diagrams obtained using the average of maximum persistence of 10 realizations of the stochastic Eq. (3.7). The third row shows the same stability diagrams using the maximum persistence of the pointwise average of 20 realizations. The overlaid white line in rows two and three is the same as the black line in the noise-free stability diagram. Note that the color scale is different for each row. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

realizations which are out of phase (e.g. see the last example time series of Fig. 2). As discussed in Section 2.3, signals with a small amplitude have smaller maximum persistence, thus lowering the overall maximum persistence values.

For the case of noise-free turning described by Eq. (3.3), Fig. 8 (top right) shows that the persistence diagram captures the stability behavior predicted numerically using the spectral element method. This can be seen by comparing the coloring of

the figure with the stability boundary obtained using the spectral element approach represented by the overlaid white line. The figure shows that the maximum persistence switches from almost zero in the stable region to positive values over a boundary that closely follows the numerically calculated stability boundary. Therefore, even though the simulation was terminated as soon as the tool jumped out of the cut, persistence was still able to detect the onset of chatter associated with going from a stable equilibrium to an unstable oscillatory behavior via Hopf bifurcation.

For the stochastic cases, particularly the case $\delta = 0.01$, the first graph in the second and third rows of Fig. 8 show a stability behavior similar to the noise-free case for regions *a*, *b*, *c*, and *d* of the parameter space that are marked on top left graph of Fig. 8. Regions *b* and *d* appear to be most resilient to the injection of noise since the sharp increase near these locations is still particularly pronounced, even for the $\delta = 0.03$ and $\delta = 0.05$ cases. In contrast to the noise-free case, we notice that the unstable region permeates the stability boundary in the *e* region for all levels of noise (see the second and third rows of Fig. 8). The infiltration of the unstable region through boundary *e* progressively increases as the noise level is increased. This can be inferred from the higher persistence values in that region which suggests the emergence of oscillatory behavior indicative of chatter. This shift in the stability boundary due to the influence of stochastic parameters can be attributed to stochastic resonance [2].

A careful investigation of regions *a*, *b*, *c*, and *d* in rows 2 and 3 of Fig. 8 shows that as δ increases the maximum persistence in regions *a* and *c* decreases while it increases in regions *b* and *d*. This suggests that noise can have a stabilizing effect near the *a* and *b* regions of the stability boundary.

Finally, if we compare rows 2 (average of maximum persistence) and 3 (maximum persistence of the average) of Fig. 8, we notice that the average of maximum persistence retains and represents more information than the maximum persistence of the average. Specifically, the regions where stability switches as indicated by a jump in the maximum persistence are more well-defined in row 2. Further, the changes in the maximum persistence values are smoother in row 2 than in row 3, which can be seen by the more pronounced boundaries for regions with close maximum persistence values. This can be explained by (1) the stability of persistence diagrams to noise and (2) by the invariance of persistence to translation and rotation of the point cloud.

To elaborate on (2), we note that multiple realizations of continuous time series from a stochastic process typically include variability in rate and amplitude. This means that the resulting point clouds of the different realizations will be translated and rotated relative to each other. Consequently, directly taking the pointwise average of the replicates, as was done in row 3 of Fig. 8, destroys some of the data. An example of this can also be seen from the time series of the replicates in the third column of Fig. 6 where the shift in the signals due to stochasticity causes part of the information in the time series to be lost when the average is taken. Persistence is invariant to translation and rotation of the point cloud; therefore, taking the average of maximum persistence preserves more information from the different realizations since persistence in some sense 'aligns' the salient topological features of the different realizations before taking the average of their maximum persistence.

7. Conclusions

The results of this study show that persistent homology captures the change in the qualitative behavior of the system associated with a bifurcation. In the context of machining dynamics, this means that persistence is a viable tool for detecting the onset of chatter.

The approach we describe is general and applicable to a variety of dynamical systems. However, in the context of chatter detection, one of the salient features of our approach in comparison to existing chatter detection methods is that it uses the topology of the underlying manifold of the dynamical system in order to detect behavioral changes indicative of chatter. Further, the current approach allows for the investigation of stochastic models with nonlinearities and delays in a systematic and quantitative way.

Although the datasets that we investigated were generated using simulation, the developed method is equally applicable to experimental data. The analysis showed that increasing the noise intensity changes the side of the stability lobes that is likely to experience more abrupt transitions to chatter. Further, in addition to the destabilizing effect that noise can have due to the stochastic resonance mechanism, it can also have a stabilizing effect as evidenced by the decrease of maximum persistence when increasing δ in regions *a* and *c* of Fig. 8. The stabilizing/destabilizing effect can be explained by the coexistence of attractors in the process with possible stochastic resonance. Future work includes fine-tuning the method so that the stability boundary can be identified with higher resolution especially at higher noise levels. The results in Fig. 8 also show that when multiple realizations of the stochastic process are available, then combining the information that they contain by utilizing the resulting persistence diagrams instead of directly using the replicate time series leads to better results. This is largely due to the invariance of persistence diagrams to translation and rotation of the point cloud and their resistance to noise.

In the future, we are interested in using different pipelines for the data with methods that more closely reflect the qualities of the system of study. The approach can be improved by using the extended Takens theorem for stochastic systems [68–70] especially in the context of delay models. We hypothesize that the loss of stability structure for greater amounts of noise in region *e* of Fig. 8 is due to the inadequacy of Takens standard theorem for highly stochastic equations. We remark, however, that the structure observed for small amounts of noise, particularly around regions *a*–*d*, demonstrates

the ability and the potential of the described approach to detect the underlying dynamics even with poor embedding of the point cloud.

In addition, this research has many further questions in the area of persistence. First, we are using only a small sliver of the information stored in the persistence diagrams, largely for ease of computation and visualization. It would be fascinating to see if the addition of the full information from the persistence diagram, or the addition of the newly developed statistical models for persistence diagrams [87–92] would improve the analysis. In addition, we want to find fast and efficient ways to compute the persistent homology of the Takens embedding so that these methods can be used online. This may require applying tools such as zigzag persistence [74] to be able to add and remove simplices in order to maintain the persistence of the Takens embedding for a short history without needing to recompute it at each time. Additionally, an implementation using persistent cohomology could give more information back about the periodicity observed using the persistence diagram [66]. The task could also be aided by utilizing the newly emerging theory of multidimensional persistence [93] in order to study the homology of a changing point cloud while still taking the diameter parameter into account. These are all subjects of our current investigations.

Acknowledgments

The authors thank Liz Bradley, Jose Perea, and Jesse Berwald for extremely helpful discussions. In addition, the authors would like to thank the two anonymous reviewers for their insightful comments which greatly improved the paper. The work described in this paper is a result of a collaboration made possible by the Institute for Mathematics and Its Applications while EM was a postdoctoral fellow during the annual program on Scientific and Engineering Applications of Algebraic Topology. FAK gratefully acknowledges the support and hospitality provided by the IMA during his visits which took place in February and May of 2014, as well as the hospitality provided by the University at Albany during the summer of 2015.

References

[1] G. Quintana, J. Ciurana, Chatter in machining processes: a review, Int. J. Mach. Tools Manuf. 51 (5) (2011) 363–376, http://dx.doi.org/10.1016/j.ijmachtools.2011.01.001.
[2] R. Kuske, Multiple-scales approximation of a coherence resonance route to chatter, Comput. Sci. Eng. 8 (3) (2006) 35–43, http://dx.doi.org/10.1109/MCSE.2006.44.
[3] E. Buckwar, R. Kuske, B. L'Esperance, T. Soo, Noise-sensitivity in machine tool vibrations, Int. J. Bifurc. Chaos 16 (08) (2006) 2407–2416, http://dx.doi.org/10.1142/S021812740601615X.
[4] B. Klamecki, Enhancement of the low-level components of milling vibration signals by stochastic resonance, Proc. Inst. Mech. Eng. Part E: J. Process Mech. Eng. 218 (1) (2004) 33–41, http://dx.doi.org/10.1243/095440804322860627.
[5] A. Yilmaz, E. AL-Regib, J. Ni, Machine tool chatter suppression by multi-level random spindle speed variation, J. Manuf. Sci. Eng. 124 (2) (2002) 208–216, http://dx.doi.org/10.1115/1.1378794.
[6] J. Gradišek, E. Govekar, I. Grabec, Qualitative and quantitative analysis of stochastic processes based on measured data, II: applications to experimental data, J. Sound Vib. 252 (3) (2002) 563–572, http://dx.doi.org/10.1006/jsvi.2001.4051. ⟨http://www.sciencedirect.com/science/article/pii/S0022460X01940519⟩.
[7] J. Gradisek, I. Grabec, S. Siegert, R. Friedrich, Stochastic dynamics of metal cutting: bifurcation phenomena in turning, Mech. Syst. Signal Process. 16 (5) (2002) 831–840, http://dx.doi.org/10.1006/mssp.2001.1403.
[8] M. Barrio, K. Burrage, A. Leier, T. Tian, Oscillatory regulation of hes1: discrete stochastic delay modeling and simulation, PLoS Comput. Biol. 2 (9) (2006) e117, http://dx.doi.org/10.1371/journal.pcbi.0020117.
[9] T. Tian, K. Burrage, P.M. Burrage, M. Carletti, Stochastic delay differential equations for genetic regulatory networks, J. Comput. Appl. Math. 205 (2) (2007) 696–707, http://dx.doi.org/10.1016/j.cam.2006.02.063. special issue on evolutionary problems.
[10] M.C. Mackey, I.G. Nechaeva, Solution moment stability in stochastic differential delay equations, Phys. Rev. E 52 (4) (1995) 3366–3376, http://dx.doi.org/10.1103/PhysRevE.52.3366.
[11] S. Guillouzic, I. L'Heureux, A. Longtin, Small delay approximation of stochastic delay differential equations, Phys. Rev. E 59 (1999) 3970–3982, http://dx.doi.org/10.1103/PhysRevE.59.3970.
[12] O. Elbeyli, J. Sun, G. Unal, A semi-discretization method for delayed stochastic systems, Commun. Nonlinear Sci. Numer. Simul. 10 (1) (2005) 85–94, http://dx.doi.org/10.1016/S1007-5704(03)00095-9.
[13] T.D. Frank, P.J. Beek, Stationary solutions of linear stochastic delay differential equations: applications to biological systems, Phys. Rev. E 64 (2001) 021917, http://dx.doi.org/10.1103/PhysRevE.64.021917.
[14] E. Buckwar, Introduction to the numerical analysis of stochastic delay differential equations, J. Comput. Appl. Math. 125 (12) (2000) 297–307, http://dx.doi.org/10.1016/S0377-0427(00)00475-1. Numerical Analysis 2000. vol. VI: Ordinary differential equations and integral equations.
[15] X. Mao, S. Sabanis, Numerical solutions of stochastic differential delay equations under local lipschitz condition, J. Comput. Appl. Math. 151 (1) (2003) 215–227, http://dx.doi.org/10.1016/S0377-0427(02)00750-1.
[16] C. Yuan, W. Glover, Approximate solutions of stochastic differential delay equations with markovian switching, J. Comput. Appl. Math. 194 (2) (2006) 207–226, http://dx.doi.org/10.1016/j.cam.2005.07.004.
[17] E. Buckwar, R. Winkler, Multi-step Maruyama methods for stochastic delay differential equations, Stoch. Anal. Appl. 25 (5) (2007) 933–959, http://dx.doi.org/10.1080/07362990701540311.
[18] Y. Hu, S.-E. A. Mohammed, F. Yan, Discrete-time approximations of stochastic delay equations: the Milstein scheme, Ann. Probab. 32 (1A) (2004) 265–314, http://dx.doi.org/10.1214/aop/1078415836.
[19] C.T.H. Baker, E. Buckwar, Exponential stability in p-th mean of solutions, and of convergent euler-type solutions, of stochastic delay differential equations, J. Comput. Appl. Math. 184 (2) (2005) 404–427, http://dx.doi.org/10.1016/j.cam.2005.01.018.
[20] J. Luo, A note on exponential stability in pth mean of solutions of stochastic delay differential equations, J. Comput. Appl. Math. 198 (1) (2007) 143–148 , http://dx.doi.org/10.1016/j.cam.2005.11.019.
[21] X. Mao, Exponential stability of equidistant Euler–Maruyama approximations of stochastic differential delay equations, J. Comput. Appl. Math. 200 (1) (2007) 297–316, http://dx.doi.org/10.1016/j.cam.2005.11.035.

[22] I. Jolliffe, Principal Component Analysis, John Wiley & Sons, Ltd, New York, NY, 2005. http://dx.doi.org/10.1002/0470013192.bsa501.
[23] H. Abdi, L.J. Williams, Principal component analysis, Wiley Interdiscip. Rev.: Comput. Stat. 2 (4) (2010) 433–459, http://dx.doi.org/10.1002/wics.101.
[24] I. Borg, P.J.F. Groenen, Modern Multidimensional Scaling: Theory and Applications, Springer, New York, NY, 2005.
[25] S. Roweis, L. Saul, Nonlinear dimensionality reduction by locally linear embedding, Science 290 (5500) (2000) 2323–2326, http://dx.doi.org/10.1126/science.290.5500.2323.
[26] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, Neural Comput. 15 (6) (2003) 1373–1396, http://dx.doi.org/10.1162/089976603321780317.
[27] D.L. Donoho, C. Grimes, Hessian eigenmaps: locally linear embedding techniques for high-dimensional data, Proc. Natl. Acad. Sci. 100 (10) (2003) 5591–5596, http://dx.doi.org/10.1073/pnas.1031596100.
[28] Z. Zhang, H. Zha, Principal manifolds and nonlinear dimension reduction via local tangent space alignment, SIAM J. Sci. Comput. 26 (2002) 313–338.
[29] R.R. Coifman, S. Lafon, Diffusion maps, Appl. Comput. Harmon. Anal. 21 (1) (2006) 5–30, http://dx.doi.org/10.1016/j.acha.2006.04.006. Special Issue: Diffusion Maps and Wavelets.
[30] B. Nadler, S. Lafon, R.R. Coifman, I.G. Kevrekidis, Diffusion maps, spectral clustering and reaction coordinates of dynamical systems, Appl. Comput. Harmon. Anal. 21 (1) (2006) 113–127, http://dx.doi.org/10.1016/j.acha.2005.07.004. Special Issue: Diffusion Maps and Wavelets.
[31] A. Singer, R.R. Coifman, Non-linear independent component analysis with diffusion maps, Appl. Comput. Harmon. Anal. 25 (2) (2008) 226–239, http://dx.doi.org/10.1016/j.acha.2007.11.001.
[32] T. McKelvey, G. Guerin, Non-parametric frequency response estimation using a local rational model, in: 16th IFAC Symposium on System Identification The International Federation of Automatic Control Brussels, Belgium, July 11–13, vol. 16, 2012, pp. 49–54. http://dx.doi.org/10.3182/20120711-3-BE-2027.00299.
[33] H. Kantz, T. Schreiber, Nonlinear Time Series Analysis, Cambridge University Press, New York, NY, 2004.
[34] F.C. Moon, T. Kalmar-Nagy, Nonlinear models for complex dynamics in cutting materials, Philos. Trans. R. Soc. London Ser. A: Math. Phys. Eng. Sci. 359 (1781) (2001) 695–711, http://dx.doi.org/10.1098/rsta.2000.0751.
[35] L. Fontanella, M. Granturco, Spectral analysis in frequency and time domain for noisy time series, in: H.-H. Bock, M. Chiodi, A. Mineo (Eds.), Advances in Multivariate Data Analysis, Studies in Classification, Data Analysis, and Knowledge Organization, Springer, Berlin, Heidelberg, 2004, pp. 67–79, http://dx.doi.org/10.1007/978-3-642-17111-6_6.
[36] Z. Yao, D. Mei, Z. Chen, On-line chatter detection and identification based on wavelet and support vector machine, J. Mater. Process. Technol. 210 (5) (2010) 713–719, http://dx.doi.org/10.1016/j.jmatprotec.2009.11.007.
[37] H. Edelsbrunner, D. Letscher, A. Zomorodian, Topological persistence and simplification, in: 41st Annual Symposium on Foundations of Computer Science, 2000. Proceedings, 2000, pp. 454–463. http://dx.doi.org/10.1109/SFCS.2000.892133.
[38] A. Zomorodian, G. Carlsson, Computing persistent homology, Discrete Comput. Geom. 33 (2) (2004) 249–274, http://dx.doi.org/10.1007/s00454-004-1146-y.
[39] H. Edelsbrunner, J. Harer, Computational Topology: An Introduction, American Mathematical Society, Providence, RI, USA, 2010.
[40] G. Carlsson, Topological pattern recognition for point cloud data, Acta Numer. 23 (2014) 268–289, http://dx.doi.org/10.1017/S0962492914000051.survey.
[41] J. Brown, T. Gedeon, Structure of the afferent terminals in terminal ganglion of a cricket and persistent homology, PLoS ONE 7 (5) (2012) e37278, http://dx.doi.org/10.1371/journal.pone.0037278.
[42] Y. Dabaghian, F. Mémoli, L. Frank, G. Carlsson, A topological paradigm for hippocampal spatial map formation using persistent homology, PLoS Comput. Biol. 8 (8) (2012) e1002581, http://dx.doi.org/10.1371/journal.pcbi.1002581.
[43] Chan, J.M., Carlsson, G., Rabadan, R., 2013. Topology of viral evolution. Proceedings of the National Academy of Sciences 110 (46), 18566-18571, http://dx.doi.org/10.1073/pnas.1313480110.
[44] P.G. Camara, A.J. Levine, R. Rabadan, Inference of ancestral recombination graphs through topological data analysis. http://arxiv.org/abs/1505.05815, May 2015.
[45] D. Taylor, F. Klimm, H. A. Harrington, M. Kramár, K. Mischaikow, M.A. Porter, P.J. Mucha, Topological data analysis of contagion maps for examining spreading processes on networks, Nature Commun. 6 (7723). http://dx.doi.org/10.1038/ncomms8723.
[46] P. Bendich, S. Chin, J. Clarke, J. deSena, J. Harer, E. Munch, A. Newman, D. Porter, D. Rouse, N. Strawn, A. Watkins, Topological and statistical behavior classifiers for tracking applications. http://arxiv.org/abs/1406.0214, June 2014.
[47] D. Rouse, A. Watkins, D. Porter, J. Harer, P. Bendich, N. Strawn, E. Munch, J. DeSena, J. Clarke, J. Gilbert, P. Chin, A. Newman, Feature-aided multiple hypothesis tracking using topological and statistical behavior classifiers, in: Proceedings of SPIE, vol. 9474, 2015, pp. 94740L–94740L–12. http://dx.doi.org/10.1117/12.2179555.
[48] L. Sjoberg, K. Knudson, Theoretical geometry, critical theory, and concept spaces in $\mathbb{R}$. http://arxiv.org/abs/1506.01104, June 2015.
[49] M. Ahmed, B.T. Fasy, C. Wenk, Local persistent homology based distance between maps, in: Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM, New York, NY, USA, 2014, pp. 43–52.
[50] V. de Silva, R. Ghrist, Coverage in sensor networks via persistent homology, Algebraic Geom. Topol. 7 (2007) 339–358, http://dx.doi.org/10.2140/agt.2007.7.339.
[51] E. Munch, M. Shapiro, J. Harer, Failure filtrations for fenced sensor networks, Int. J. Robot. Res. 31 (9) (2012) 1044–1056, http://dx.doi.org/10.1177/0278364912451671.
[52] G. Carlsson, T. Ishkhanov, V. de Silva, A. Zomorodian, On the local behavior of spaces of natural images, Int. J. Comput. Vis. 76 (2008) 1–12, http://dx.doi.org/10.1007/s11263-007-0056-x.
[53] F. Takens, Detecting strange attractors in turbulence, in: D. Rand, L.-S. Young (Eds.), Dynamical Systems and Turbulence, Warwick 1980, Lecture Notes in Mathematics, vol. 898, Springer, Berlin, Heidelberg, 1981, pp. 366–381.http://dx.doi.org/10.1007/BFb0091924.
[54] V. de Silva, P. Skraba, M. Vejdemo-Johansson, Topological analysis of recurrent systems, in: NIPS 2012 Workshop on Algebraic Topology and Machine Learning, 2012.
[55] J. Berwald, M. Gidea, M. Vejdemo-Johansson, Automatic recognition and tagging of topologically different regimes in dynamical systems. arXiv:1312.2482.
[56] K.A. Brown, K.P. Knudson, Nonlinear statistics of human speech data, Int. J. Bifurc. Chaos 19 (07) (2009) 2307–2319, http://dx.doi.org/10.1142/S0218127409024086.
[57] S. Emrani, T. Gentimis, H. Krim, Persistent homology of delay embeddings and its application to wheeze detection, IEEE Signal Process. Lett. 21 (4) (2014) 459–463, http://dx.doi.org/10.1109/LSP.2014.2305700.
[58] J. Perea, J. Harer, Sliding windows and persistence: an application of topological methods to signal analysis, Found. Comput. Math. 15 (3) (2015) 799–838, http://dx.doi.org/10.1007/s10208-014-9206-z.
[59] J. A. Perea, A. Deckard, J. Harer, S. Haase, Sw1pers: sliding windows and 1-persistence scoring; discovering periodicity in gene expression time series data, Preprint.
[60] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Y. Mileyko, Lipschitz functions have l_p-stable persistence, Found. Comput. Math. 10 (2) (2010) 127–139, http://dx.doi.org/10.1007/s10208-010-9060-6.
[61] A. Deckard, R.C. Anafi, J.B. Hogenesch, S.B. Haase, J. Harer, Design and analysis of large-scale biological rhythm studies: a comparison of algorithms for detecting periodic signals in biological data, Bioinformatics 29 (24) (2013) 3174–3180, http://dx.doi.org/10.1093/bioinformatics/btt541.
[62] Z. Alexander, J.D. Meiss, E. Bradley, J. Garland, Iterated function system models in data analysis: detection and separation, Chaos: Interdiscip. J. Nonlinear Sci. 22 (2) (2012). http://dx.doi.org/10.1063/1.3701728.x.

[63] Z. Alexander, A topology-based approach for nonlinear time series with applications in computer performance analysis (Ph.D. thesis), University of Colorado at Boulder, 2012.
[64] M. Vejdemo-Johansson, F. Pokorny, P. Skraba, D. Kragic, Cohomological learning of periodic motion, Appl. Algebra Eng. Commun. Comput. (2015) 1–22, http://dx.doi.org/10.1007/s00200-015-0251-x.
[65] M. Robinson, Topological Signal Processing, Springer, Berlin, Germany, 2014.
[66] V. de Silva, D. Morozov, M. Vejdemo-Johansson, Persistent cohomology and circular coordinates, Discrete Comput. Geom. 45 (4) (2011) 737–759, http://dx.doi.org/10.1007/s00454-011-9344-x.
[67] J. Stark, D. Broomhead, M. Davies, J. Huke, Takens embedding theorems for forced and stochastic systems, Nonlinear Anal.: Theory Methods Appl. 30 (8) (1997) 5303–5314, in: Proceedings of the Second World Congress of Nonlinear Analysts, doi: 10.1016/ S0362-546X(96)00149-6.
[68] M.R. Muldoon, D.S. Broomhead, J.P. Huke, R. Hegger, Delay embedding in the presence of dynamical noise, Dyn. Stab. Syst. 13 (2) (1998) 175–186, http://dx.doi.org/10.1080/02681119808806259.
[69] J. Stark, D. Broomhead, M. Davies, J. Huke, Delay embeddings for forced systems. II. Stochastic forcing, J. Nonlinear Sci. 13 (6) (2003) 519–577, http://dx.doi.org/10.1007/s00332-003-0534-4.
[70] J.C. Robinson, A topological delay embedding theorem for infinite-dimensional dynamical systems, Nonlinearity 18 (5) (2005) 2–35.
[71] J.R. Munkres, Elements of Algebraic Topology, Addison Wesley, Cambridge, MA, USA, 1993.
[72] A. Hatcher, Algebraic Topology, Cambridge University Press, Cambridge, MA, USA, 2002.
[73] R. Ghrist, Barcodes: the persistent topology of data, Bull. Am. Math. Soc. 45 (2008) 61–75. survey.
[74] G. Carlsson, V. de Silva, Zigzag persistence, Found. Comput. Math. 10 (4) (2010) 367–405, http://dx.doi.org/10.1007/s10208-010-9066-0.
[75] D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, Discrete Comput. Geom. 37 (1) (2007) 103–120, http://dx.doi.org/10.1007/s00454-006-1276-5.
[76] J. Tlusty, Manufacturing Processes and Equipment, 1st edition, Prentice Hall, Upper Saddle River, NJ, 2000.
[77] G. Stepan, Modeling nonlinear regenerative effects in metal cutting, Philos. Trans. R. Soc. London Ser. A: Math. Phys. Eng. Sci. 359 (1781) (2001) 739–757, http://dx.doi.org/10.1098/rsta.2000.0753.
[78] T. Insperger, D.A. Barton, G. Stepan, Criticality of hopf bifurcation in state-dependent delay model of turning processes, Int. J. Non-Linear Mech. 43 (2) (2008) 140–149, http://dx.doi.org/10.1016/j.ijnonlinmec.2007.11.002.
[79] B. Oksendal, Stochastic Differential Equations, 6th edition, Springer, Berlin, Germany, 2007.
[80] F.A. Khasawneh, B.P. Mann, A spectral element approach for the stability of delay systems, Int. J. Numer. Methods Eng. 87 (6) (2011) 566–592, http://dx.doi.org/10.1002/nme.3122.
[81] Y. Altintaş, E. Budak, Analytical prediction of stability lobes in milling, CIRP Ann. 44 (1) (1995) 357–362.
[82] T. Insperger, G. Stépán, Updated semi-discretization method for periodic delay-differential equations with discrete delay, Int. J. Numer. Methods 61 (2004) 117–141.
[83] E.A. Butcher, O.A. Bobrenkov, On the Chebyshev spectral continuous time approximation for constant and periodic delay differential equations, Commun. Nonlinear Sci. Numer. Simul. 16 (3) (2011) 1541–1554, http://dx.doi.org/10.1016/j.cnsns.2010.05.037.
[84] D. Higham, An algorithmic introduction to numerical simulation of stochastic differential equations, SIAM Rev. 43 (3) (2001) 525–546, http://dx.doi.org/10.1137/S0036144500378302.
[85] M.B. Kennel, R. Brown, H.D.I. Abarbanel, Determining embedding dimension for phase-space reconstruction using a geometrical construction, Phys. Rev. A 45 (1992) 3403–3411, http://dx.doi.org/10.1103/PhysRevA.45.3403.
[86] J. Harer, J. Slaczedek, P. Bendich, Ripscollapse: discrete morse theory and fast computation of one-dimensional persistence, Manuscript, Duke University.
[87] P. Bubenik, Statistical topological data analysis using persistence landscapes, J. Mach. Learn. Res. 16 (2015) 77–102.
[88] B.T. Fasy, F. Lecci, A. Rinaldo, L. Wasserman, S. Balakrishnan, A. Singh, Confidence sets for persistence diagrams, Ann. Stat. 42 (6) (2014) 2301–2339, http://dx.doi.org/10.1214/14-AOS1252.
[89] Y. Mileyko, S. Mukherjee, J. Harer, Probability measures on the space of persistence diagrams, Inverse Probl. 27 (12) (2011) 124007.
[90] K. Turner, Y. Mileyko, S. Mukherjee, J. Harer, Fréchet means for distributions of persistence diagrams, Discrete Comput. Geom. 52 (1) (2014) 44–70, http://dx.doi.org/10.1007/s00454-014-9604-7.
[91] Munch, E., Turner, K., Bendich, P., Mukherjee, S., Mattingly, J., Harer, J., 2015. Probabilistic Fréchet means for time varying persistence diagrams. Electron. J. Statist. 9 (1), 1173-1204, http://dx.doi.org/10.1214/15-EJS1030.
[92] Blumberg, A.J., Gal, I., Mandell, M.A., Pancia, M., 2014. Robust Statistics, Hypothesis Testing, and Confidence Intervals for Persistent Homology on Metric Measure Spaces. Foundations of Computational Mathematics 14 (4), 745-789, http://dx.doi.org/10.1007/s10208-014-9201-4.
[93] G. Carlsson, A. Zomorodian, The theory of multidimensional persistence, Discrete Comput. Geom. 42 (1) (2009) 71–93, http://dx.doi.org/10.1007/s00454-009-9176-0.