MANCHESTER
1824

The University of Manchester

Smoothing non-smooth systems with low-pass filters

Hook, JL

2013

MIMS EPrint: **2013.15**

ISSN 1749-9097

Smoothing non-smooth systems with low-pass filters

James Hook

Alan Turing Building, Oxford Rd, Manchester, UK

Abstract

Low pass filters, which are used to remove high frequency noise from time series data, smooth the signals they are applied to. In this paper we examine the action of low pass filters on discontinuous or non-differentiable signals from non-smooth dynamical systems. We show that the application of the filter is equivalent to a change of variables, which transforms the non-smooth system into a smooth one. We examine this smoothing action on a variety of examples and demonstrate how it is useful in the calculation of a non-smooth system's Lyapunov spectrum.

Introduction

Non-smooth dynamical systems are used to model mechanical systems with impacts or friction, as well as control systems with switching between distinct modes of operation. Non-smooth systems are also interesting mathematically as they generically exhibit bifurcation structures that would be impossible or of high co-dimension in the space of smooth systems [1].

In this paper we introduce the notion of smoothing a non-smooth system with a low-pass filter. The idea is that the filter's action on the time-series can be used to construct a change of variables that actually transforms a non-smooth system into a smooth one. There are some subtleties here, the 'smoothed' system will not be smooth everywhere as singular discontinuities (grazes and chattering points) will be mapped to singularities in the new flow, also the transformation and the smoothed system will typically be impossible to compute analytically. However we will still be able to obtain it for simple examples or numerically for more complex systems. To apply a smoothing transformation numerically to an orbit we simply apply the associated low pass filter to the time-series. Indeed whenever an engineer analyses data from

Preprint submitted to Physica D *July 16, 2013*

a non-smooth system that has been filtered they are inadvertently studying a 'smoothed' system of the sort presented here.

There are therefore two complimentary reasons for trying to undersand the action of these smoothing transformations. Firstly we might find the smoothing action useful or interesting in its own right (we will show that it is useful for computing Lyapunov exponents) and secondly such systems are already being investigated whenever experimental data is smoothed with a low pass filter.

This paper is organised as follows. In section 1 we show how a non-smooth system can be transformed into a smooth one using a change of variables. The approach used in section 1 is rather ac hoc and does not use low pass filters but it allows us to understand the link between the smooth and non-smooth system in as simple a setting as possible. In section 2 we show how a similar smoothing action can be achieved using our low-pass filter formulation. We examine some analytic examples and consider some of the signal processing issues associated with the transformation. In section 3 we briefly explain the state space reconstruction method which enables us to model a differentiable dynamical system from its time-series data and examine some simple numerical examples. In section 4 we argue that linear stability (when it exists) is preserved by the smoothing procedure. In section 5 we apply the smoothing procedure to time series from a impacting Duffing oscillator and calculate its Lyapunov spectrum using a time-series method that relies on differentiability.

1. Ad hoc smoothing

Consider a mass on a linear spring whose motion is obstructed by a wall placed at the spring's natural length. Suppose that when the mass hits the wall it bounces off it elastically with coefficient of restitution c. Let $x(t)$ measure the distance from the wall to the mass. The motion of the mass is governed by $\ddot{x} = -x$, along with the rule that whenever $\lim_{\tau \to t} x(\tau) = 0$ we set $\dot{x}(t) = \lim_{\tau \to t} -c\dot{x}(\tau)$.

Orbits to this system are composed of a series of smaller and smaller semi-circles, see figure 1. A solution evolves by describing one of the semi-circles until it reaches $x = 0$, when it instantaneously jumps to the start of a new smaller semi-circle and so on. This roughly periodic behaviour is just like that of a smooth system with a stable fixed point, where solutions spiral into the equilibrium.

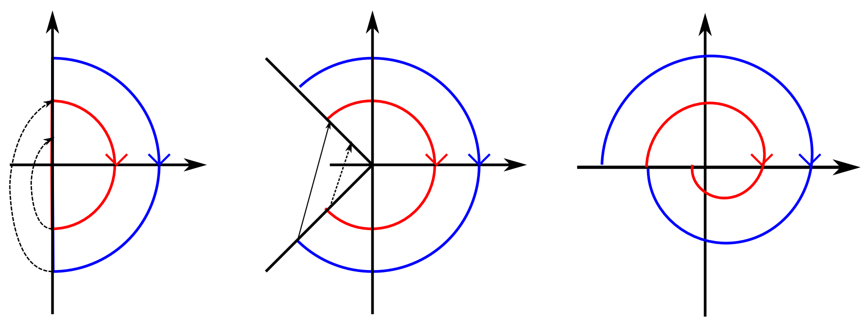

Figure 1: Left to right: deforming the discontinuous system into a smooth system.

Indeed we can imagine sticking a pin into the origin of this picture and stretching the space around it to fill the plane. The two sides of the boundary would meet and it would be possible to glue them together so that the jump 'take offs' and 'landings' joined together. If any kinks in the picture could then be ironed out we would have something that looked exactly like the stable equilibrium of a smooth system.

It turns out that for this simple example we can formulate a transformation that has these exact properties. In polar co-ordinates

$$T\begin{pmatrix} r \\ \theta \end{pmatrix} = \begin{pmatrix} rc^{\frac{\theta}{\pi}} \\ 2\theta \end{pmatrix},$$

maps the semi-circle starting at $(\dot{x}, 0)$ to the 360 degree spiral starting at $(\dot{x}, 0)$ and finishing at $(c\dot{x}, 0)$. T transforms the original non-smooth system to a smooth system governed by

$$\frac{d}{dt}\begin{bmatrix} p \\ q \end{bmatrix} = \begin{bmatrix} \frac{\log(c)}{\pi} & 1 \\ -1 & \frac{\log(c)}{\pi} \end{bmatrix}\begin{bmatrix} p \\ q \end{bmatrix}.$$

Let $\phi : X \times \mathbb{R}_+ \mapsto X$ be the flow of the original discontinuous system and $\varphi : \mathbb{R}^2 \times \mathbb{R}_+ \mapsto \mathbb{R}^2$ be the flow of the new smoothed system. The transformation T provides the commutation

$$\phi_t(x, \dot{x}) = T^{-1} \circ \varphi_t \circ T(x, \dot{x}),$$

so that we can substitute one flow for another. Likewise their stability is related by

$$\frac{d\phi_t(x', \dot{x}')}{d(x', \dot{x}')}|_{(x,\dot{x})} = \frac{dT^{-1}(p,q)}{d(p,q)}|_{\varphi_t \circ T(x,\dot{x})} \times \frac{d\varphi_t(p,q)}{d(p,q)}|_{T(x,\dot{x})} \times \frac{dT(x', \dot{x}')}{d(x', \dot{x}')}|_{(x,\dot{x})}.$$

This alternative expression is much simpler to evaluate as we no longer have to worry about repeated application of saltation matrices every time the orbit crosses the discontinuity. Instead we only need to evaluate the stability of the smooth flow then multiply it by the derivatives of T. Moreover since the derivatives of T and its inverse are everywhere bounded the Lupanov spectrums of the two systems are identical. It is easy to show from the smoothed system that both Lupanov exponents are equal to $\log(c)/\pi$.

We are able to play the same game with the bouncing ball system. This evolves according to $\ddot{x} = -g$ along with the rule that whenever $\lim_{\tau\to t} x(\tau) = 0$, we set $\dot{x}(t) = \lim_{\tau\to t} -c\dot{x}(\tau)$. Here the analysis is only a little more complicated, we start by mapping parabolas to 360 degree spirals to obtain a transformation between the discontinuous bouncing ball system and the smooth (everywhere except the origin) system

$$\frac{d}{dt}\begin{bmatrix} p \\ q \end{bmatrix} = \frac{2\pi g(c-1)}{\log c\sqrt{p^2+q^2}}\begin{bmatrix} \frac{\log c}{2\pi} & 1 \\ -1 & \frac{\log c}{2\pi} \end{bmatrix}\begin{bmatrix} p \\ q \end{bmatrix}.$$

It should come as no surprise that this 'smoothed' system is not quite smooth. The conjugacy between its flow and that of the bouncing ball means they must share stability properties and the bouncing ball is a singular system; all orbits reach the origin in finite time.

It would be fantastic if we could explicitly construct such conjugacies for more complicated non-smooth systems. Unfortunately although our filter based transformation has the desired smoothing action, explicitly applying it to obtain the smoothed system is not typically possible as it requires integrating solutions to the original problem. However we will show in the next section that it is still possible to examine the smoothed system by smoothing time-series data recorded from the non-smooth system.

2. Smoothing with low-pass filters

A finite impulse response filter Ψ is a linear operator given by

$$\Psi(f)(t) = \int_{-w}^{0} f(t+\tau)h(\tau)d\tau,$$

where $h(\tau)$ is the kernel and w the window. We can use filters like these to create a smoothing transformation. Let $\phi : \mathbb{R}^n \times \mathbb{R} \mapsto \mathbb{R}^n$ be the flow of a non-smooth system with state space $\mathbb{R}^n$. Now define $\Phi : \mathbb{R}^n \mapsto \mathbb{R}^n$ by

$$[\Phi(x)]_i = \int_{-w}^{0} [\phi_\tau(x)]_i h(\tau) d\tau,$$

so that $\Phi(x)$'s ith component is calculated by integrating the value of the ith co-ordinate of x's orbit up to w seconds backwards in time. For the time-being we will assume that Φ is an invertible map. Given this assumption Φ induces a new flow φ defined by

$$\varphi_t(y) = \Phi \circ \phi_t \circ \Phi^{-1}(y).$$

So that if we think of $[\phi_t(x)]_i$ as a function of time t and likewise for φ, then we have

$$\begin{array}{ccc} x & \xrightarrow{\quad \Phi \quad} & \Phi(x) \\ \phi \big\downarrow & & \big\downarrow \varphi \\ \phi(x) & \xrightarrow[\Psi]{\quad\quad} & \Psi[\phi(x)] = \varphi[\Phi(x)] \end{array}$$

Or in words, the unfiltered orbits from the transformed (smoothed) system are identical to the filtered orbits of the original system. Whenever we analyse time-series data from a non-smooth system that has been filtered (to reduce noise say) we are inadvertently studying one of these smoothed systems. It is therefore important to understand how this smoothing process affects the features that we are interested in e.g. stability and grazing points.

2.1. Invertibility

For our purposes it is essential that Φ be invertible. This requirement can be split into two parts. Firstly we require the filter Ψ to be invertible as an operator on the time series $\phi_t(x)$. Secondly given this first condition we still require that the transformation Φ itself is an injection. Both of these problems are well studied in the context of smooth systems [2, 3].

Filters, as defined in (6), are best described in Fourier space where we have

$$\widehat{\Psi(f)}(s) = \widehat{f}(s) \times H(s),$$

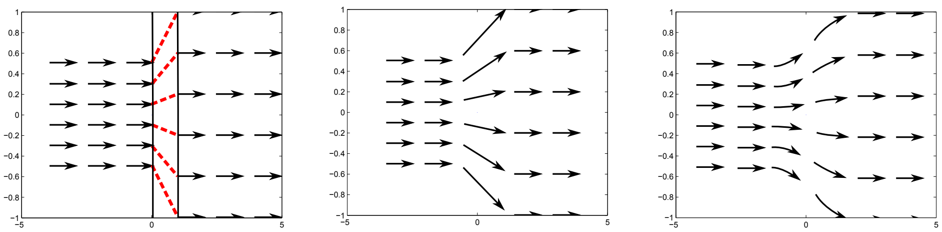

Figure 2: Left to right: vector field of original discontinuous system, once smoothed system and twice smoothed system.

where $\widehat{}$ denotes the Fourier transform and H is the transfer function

$$H(s) = \int_0^w h(-\tau) \exp[-2\pi i s \tau] d\tau.$$

If the application of a filter is to be invertible it is essential that $H(s) = 0$ only when $\widehat{f}(s) = 0$ also. Of course if we are applying a filter to remove noise the transformation will not be invertible (we can not expect to recover deleted noise) but we must still compare the spectrum of the input data to the transfer function of the filter to ensure that the only information lost is in a band consigned as noise.

As we will show in example 3 invertibility of Ψ does not guarantee invertibility of Φ. In these cases we use the method of delays to construct an invertible map with the required properties. Taken's theorem states that for generic smooth flow ϕ, delay d and smooth measuring function f, there exists finite m such that

$$F(x) = \Big(f[x], f[\phi_{-d}(x)], ..., f[\phi_{-md}(x)]\Big),$$

is invertible and provides the conjugacy between ϕ and a diffeomorphically equivalent system φ. Of course our systems are non-smooth and our measuring function depend on the flow so they are not generic. Therefore Taken's theorem gives us no guarantees but we follow its spirit and find that the method of delays works well in this non-smooth setting.

2.2. Example 1 - Moving average transformation applied twice to a discontinuous system

Let ϕ be the flow of the system governed by $(\dot{x}, \dot{y}) = (1, 0)$ along with the rule that whenever $\lim_{\tau \to t} x(\tau) = 0$, we set $(x(t), y(t)) = (1, \lim_{\tau \to t} 2y)$, see figure 2. The moving average transformation Φ is the simplest of our smoothing transformations with window $w = 1$ and constant kernel $h = 1$.

The transformation is piecewise smooth on smooth on the state space of our system. For $x < 0$ or $x \geq 2$ there are no discontinuities in the one-second backward time flow so that

$$\Phi \begin{bmatrix} x \\ y \end{bmatrix} = \int_{-1}^{0} \begin{bmatrix} x - \tau \\ y \end{bmatrix} d\tau = \begin{bmatrix} x - \frac{1}{2} \\ y \end{bmatrix}.$$

For $1 \geq x < 2$ there will be a discontinuity in the one-second backwards time flow so that the integral expression for Φ will contain contributions from before and after the jump,

$$\Phi \begin{bmatrix} x \\ y \end{bmatrix} = \int_{-1}^{1-x} \begin{bmatrix} x - 1 - \tau \\ \frac{y}{2} \end{bmatrix} d\tau + \int_{1-x}^{0} \begin{bmatrix} x - \tau \\ y \end{bmatrix} d\tau = \begin{bmatrix} 2x - \frac{5}{2} \\ \frac{xy}{2} \end{bmatrix}.$$

The transformed system φ is the governed by the non-smooth ODE

$$\begin{bmatrix} \dot{p} \\ \dot{q} \end{bmatrix} = \frac{d\Phi}{\dot{(x, y)}} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}_{\Phi^{-1}(p,q)},$$

and inverting Φ and substituting gives

$$\begin{bmatrix} \dot{p} \\ \dot{q} \end{bmatrix} = \begin{cases} \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \text{for } p < -0.5, \\ \begin{bmatrix} 2 \\ \frac{2q}{p+\frac{5}{2}} \end{bmatrix}, & \text{for } -0.5 \leq p < 1.5, \\ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \text{for } 1.5 \leq p. \end{cases}$$

Applying the moving average transformation we have obtained a non-differentiable but continuous system conjugate to our original discontinuous system.

We can now apply the transformation a second time, which is equivalent to one application of a smoother filter with window $w = 2$ and kernel

$$h(\tau) = \begin{cases} \tau + 2, & \text{for } -2 \leq \tau < -1, \\ -\tau, & \text{for } -1 \leq \tau \leq 0. \end{cases}$$

Again, this transformation is piecewise smooth but now depends on the behaviour of the two second backward time flow,

$$\Phi \begin{bmatrix} x \\ y \end{bmatrix} = \begin{cases} \begin{bmatrix} x-1 \\ y \end{bmatrix}, & \text{for } x < 0, \\ \begin{bmatrix} \frac{x^2}{2} - \frac{3}{2} \\ (\frac{(x-1)^2}{4} + \frac{1}{2})y \end{bmatrix}, & \text{for } 1 \leq x < 2, \\ \begin{bmatrix} -\frac{x^2}{2} + 4x - \frac{11}{2} \\ (1 - \frac{(x-3)^2}{4})y \end{bmatrix}, & \text{for } 2 \leq x < 3, \\ \begin{bmatrix} x-1 \\ y \end{bmatrix}, & \text{for } 3 \leq x. \end{cases}$$

The discontinuous system is transformed into the differentiable system governed by

$$\begin{bmatrix} \dot{p} \\ \dot{q} \end{bmatrix} = \begin{cases} \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \text{for } p < -1, \\ \begin{bmatrix} \sqrt{2p+3} \\ \frac{4q(\sqrt{2p+3}-1)}{(\sqrt{2p+3}-1)^2+2} \end{bmatrix}, & \text{for } -1 \leq p < 0.5, \\ \begin{bmatrix} \sqrt{5-2a} \\ \frac{2q(\sqrt{5-2a}-1)}{4-(1-\sqrt{5-2a})^2} \end{bmatrix}, & \text{for } 0.5 \leq a < 2, \\ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \text{for } 2 \leq p. \end{cases}$$

Through double application of the moving average transformation we have obtained a differentiable system conjugate to our original discontinuous system.

The non-differentiable features that would have affected the stability of the system have been smoothed out. But their effect on the dynamics has not been lost, it has been integrated into the new smooth flow. Smoothing non-smooth systems doesn't destroy information about the discontinuities, rather it encodes this information in a different way.

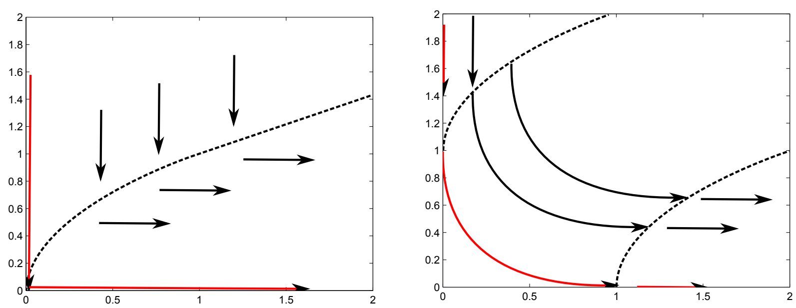

Figure 3: Left to right: vector field of original discontinuous system and once smoothed system, red curve indicates grazing orbit and its T image.

This smoothing action at regular discontinuities (jumps or switches that are not grazes or chattering points) can be shown to work in a general setting. In [4] we present normal forms for smoothing these sorts of discontinuities with the moving average transformation.

2.3. Example 2 - Moving average transformation applied once to a non-differentiable grazing system

Let ϕ be the flow governed by

$$\begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} = \begin{cases} \begin{bmatrix} 0 \\ -1 \end{bmatrix}, & \text{for } x < y^2, \\ \\ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, & \text{for } x \geq y^2. \end{cases}$$

This system has a grazing point at the origin. The moving average transformation is given by

$$\Phi \begin{bmatrix} x \\ y \end{bmatrix} = \begin{cases} \begin{bmatrix} x \\ y - \frac{1}{2} \end{bmatrix}, & \text{for } x < y^2, \\ \\ \begin{bmatrix} y^2 + \frac{(x-y^2)^2}{2} \\ y + \frac{(1-x+y^2)^2}{2} \end{bmatrix}, & \text{for } y^2 \leq x < y^2 + 1, \\ \\ \begin{bmatrix} x - \frac{1}{2} \\ y \end{bmatrix}, & \text{for } y^2 + 1 \leq x. \end{cases}$$

Unfortunately we can not explicitly invert Φ as it requires us to solve a quartic polynomial. However we can approximate the flow at the image of the graze by ignoring terms that are smaller than x or y^2 to obtain a local description of the transformed system for $(p, q) \approx (0, 0.5)$ given by

$$\begin{bmatrix} \dot{p} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} \sqrt{p} \\ \sqrt{p} \end{bmatrix} + \begin{bmatrix} 0 \\ -1 \end{bmatrix} + \begin{bmatrix} 0 & -1 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} p \\ q - \frac{1}{2} \end{bmatrix}.$$

Notice that although the square root expression will be non-differentiable at $p = 0$ the affine part of the right hand side ensures that there is no problem with non-uniqueness of solution.

It is easy to show that the transformed system is differentiable everywhere except at the image of the grazing orbit. Away from this orbit the smoothing action will be exactly as in example 1, the non-differentiable switch will be replaced with a differentiable switch and any saltation associated with the switch integrated into the new flow.

As with the regular jump and switch this smoothing action is totally general and we derive a normal form for the moving average smoothed graze in [4].

3. State space reconstructions

We have shown that filtered data from non-smooth systems resembles unfiltered data from smoothed system. Therefore in cases where we are unable to study the smoothed systems directly by explicitly computing and applying the transformation, we can instead record time-series data from the non-smooth system and filter it to obtain a time-series that is equivalent to a recording from the smoothed system.

In this section we give a very brief and informal account of the method of state space reconstruction through which one is able to construct a numerical model of a differentiable system from its time-series. For a more thorough exposition see [5].

Suppose that ϕ is a differentiable flow on a manifold $M \subset \mathbb{R}^n$. We are able to record an orbit of the system on a computer by storing the value of $\phi_t(x)$ every τ seconds. The data is a long sequence $[x(i)]_{i=1}^{N}$ with the property that $x(i+1) = \phi_\tau[x(i)]$.

Provided the stored orbit explores the manifold M sufficiently thoroughly we will be able to build a piecewise affine model of M and the time-τ map $\phi_\tau : M \mapsto M$.

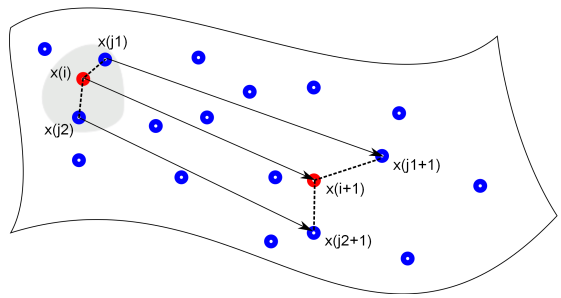

Figure 4: Reconstruction from data cartoon, grey region represents ball of radius ϵ centered at $x(i)$.

To construct an affine model of M and ϕ_τ in the vicinity of the point $x(i)$ we start by searching the time-series for other data points that are within a distance ϵ of $x(i)$. Let $x(j_1), x(j_2), ..., x(j_m)$ be the points that are close to $x(i)$. A basis for the tangent space of M at $x(i)$ is approximated by the span of the matrix

$$T(i) = \left[\begin{array}{cccc} x(j_1) - x(i) & x(j_2) - x(i) & \dots & x(j_m) - x(i) \end{array} \right],$$

which is computed by taking the SVD decomposition and ignoring basis vectors corresponding to small singular values,

$$T(i) = U(i)S(i)V(i),$$

so the basis of M's tangent space at $x(i)$ is given by $[U_1(i), U_2(i), ..., U_d(i)]$ say, where $U_k(i)$ is the kth column of $U(i)$.

The time-τ map takes points close to $x(i)$ to points close to $x(i+1)$ and we have already shown how to approximate M in these locations. Using these tangent vectors for local co-ordinates at $x(i)$ and $x(i+1)$ we have

$$\phi_\tau[x(i) + z] = x(i+1) + U(i+1)|_d \widehat{T}(i)(V^t(i)s|_d^{-1})z,$$

where

$$U(i+1)|_d = \left[\begin{array}{cccc} U_1(i+1) & U_2(i+1) & \dots & U_d(i+1) \end{array} \right],$$

$$\widehat{T}(i) = \left[\begin{array}{cccc} x(j_1+1) - x(i+1) & x(j_2+1) - x(i+1) & \dots & x(j_m+1) - x(i+1) \end{array} \right],$$

$$(V^t(i)s|_d^{-1}) = \left[\begin{array}{cccc} \frac{V_1(i)}{S_{1,1}(i)} & \frac{V_2(i)}{S_{2,2}(i)} & \dots & \frac{V_d(i)}{S_{d,d}(i)} \end{array} \right].$$

With sufficient data it is possible to reconstruct a system accurately enough to compute its Lyapunov spectrum and even predict its future behaviour.

3.1. Example 3 - moving average transformation and bump transformation applied to a noisy data from a discontinuous system

We return to the mass on a spring system used in section 1. The motion of the mass is governed by $\ddot{x} = -x$ along with the rule that whenever $\lim_{\tau \to t} x(\tau) = 0$, we set $\dot{x}(t) = \lim_{\tau \to t} -c\dot{x}(\tau)$. To set up an artificial study of noisy experimental data we simulate this system on a computer and add Gaussian noise to every variable to produce a noisy time series. The noisy data is useless for reconstructing the phase space of the system, all we see is a noisy blob, see figure 5. In order to reduce the noise in the signal we smooth it with the moving average filter. Since we are working in a numerical setting we have the time series stored as a discrete sequence so any filter will be in the form of a weighted sum rather than an integral.

The moving average smoothed data clearly has much less noise than the original signal. However if we use this data to try to reconstruct the original phase space we have two problems. Firstly, as we would expect the data is no-longer discontinuous. Smoothing the data is equivalent to smoothing the system and we have already shown that the moving average transformation will map a discontinuous system to a continuous but non-differentiable one. Secondly, the transformation is not an injection; different points in the original phase space are mapped to the same point under the smoothing, so we can not reconstruct a smooth dynamical system from this data. To remedy this problem we add a delay vector

$$D[x(t), dx(t)] = [x(t), dx(t), x(t-d)].$$

The delayed data gives a reconstruction of the state space of the continuous but non-differentiable system conjugate to the original system via the moving average transformation with delay. This smoothed system is an ODE with discontinuous right hand side whose state space is a non-differentiable cone formed from two smooth halves.

To obtain a totally smooth system we can apply a filter with smooth kernel. The bump filter has window $w = 1$ and kernel

$$h(\tau) = e^{\frac{-1}{x(x-1)}}.$$

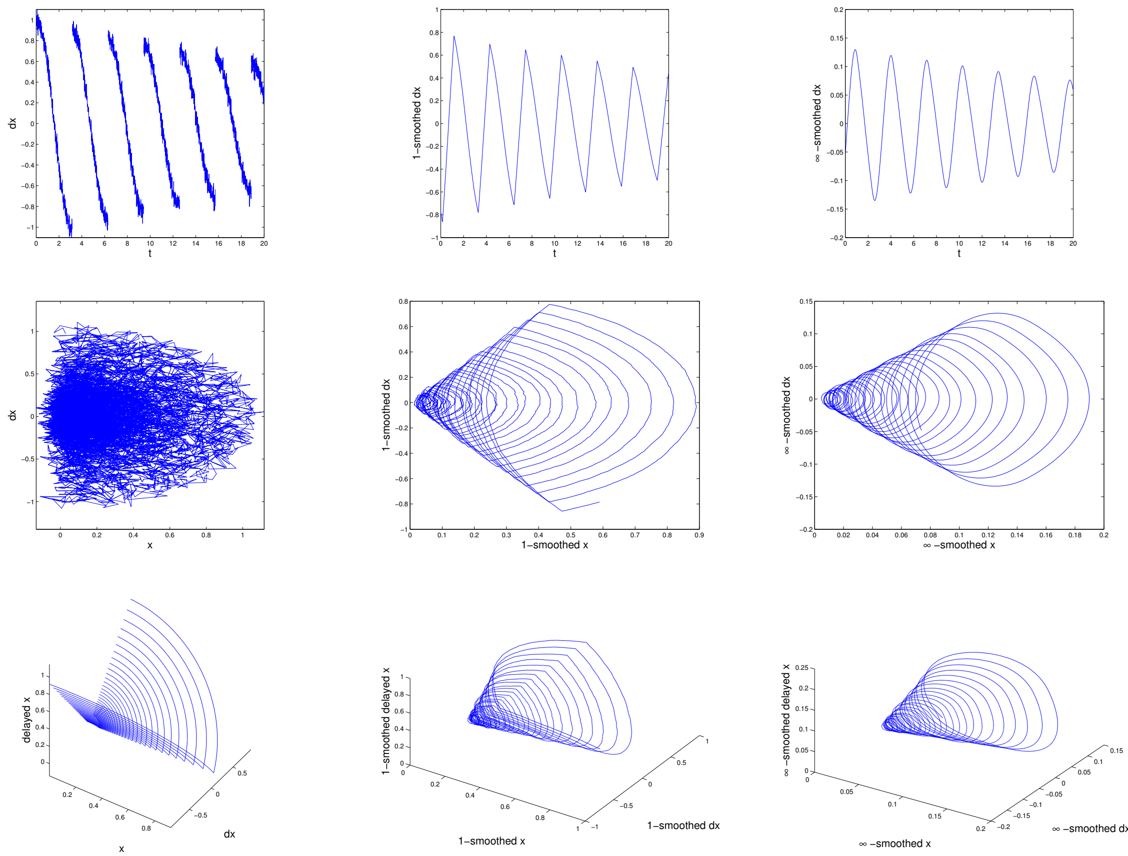

Figure 5: Left to right: non-smooth system, moving average smoothed system, bump smoothed system. Top to bottom: velocity variable versus time, phase space, delay space. (Bottom left is noise free.)

In principal this filter can completely smooth any integrable data, however since we are working numerically we instead expect to see any jumps in the first few derivatives to be removed. The exact action will depend on the stiffness of the data and the length of the filtering window. Again, the smoothed time series is not sufficient for a phase space reconstruction as the transformation is not injective, so we use the method of delays.

The smoothed delayed data gives a reconstruction of the smooth system conjugate to the original system via the bump filter with delay. This smoothed system is an ODE with smooth right hand side with state space a

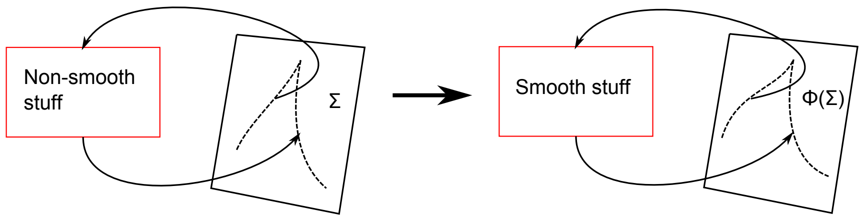

Figure 6: Poincaré map cartoon.

smooth cone. In fact if we position ourselves to 'look down' into the cone we would have a picture exactly like the smoothed system described in section 1.

We therefore find that it is impossible to apply a filter to reduce noise from a non-smooth time series without transforming the system into a smooth one. However this needn't be a problem. We argue in section 4 that the smoothed system will have the same stability properties as the original system, and we have already shown how non-smooth discontinuities are smoothed out by the transformations so that if we are interested in looking for e.g. a graze we know to look for a square root singularity in the flow.

4. Smoothness and stability

The easiest way to show that smoothing preserves stability is to consider a Poincaré return map constructed away from any discontinuities. Provided there are no discontinuities in the w-second backward time flow from the section Σ, we can be sure that Φ is smooth on Σ. The Poincaré maps of the original and smoothed systems are therefore smoothly conjugate and will have exactly the same stability properties.

An immediate consequence of this stability equivalence is that our smooth systems will often not be smooth everywhere. If a non-smooth system contains a grazing orbit, say, its return map will contain singular points which will have to be mirrored in the return map of the smoothed system.

Just as we saw in the ad hoc smoothing of the bouncing ball and the moving average smoothing of the graze discontinuity, these singular discontinuities give rise to isolated singularities in the otherwise smooth flow of the transformed system.

5. Numerical Example with stability calculation

In this section we will apply our smoothing procedure to time-series data recorded from a computer simulation of a non-linear Duffing impact oscillator. We will then use the smoothed data to calculate the Lyapunov spectrum of the system using a method that relies on the underlying system being differentiable.

The system we use is taken from [6] where Stefanski uses the coupling method to determine the largest Lyapunov exponent of the system from a numerical experiment. This provides us with a standard to test our result against. The system is governed by

$$\ddot{x} = x(1 - x^2) - 0.1\dot{x} + \cos t$$

along with the rule that whenever $\lim_{\tau \to t} x(\tau) = 0.5$, we set $\dot{x}(t) = \lim_{\tau \to t} -0.65\dot{x}(\tau)$. In order to make the system autonomous we include a forcing phase variable θ that obeys $\dot{\theta} = 1$, along with the rule that whenever $\lim_{\tau \to t} \theta(\tau) = 2\pi$, we set $\theta(t) = 0$. This autonomous formulation has two different discontinuities, one associated with the impacting in the oscillator model, and another associated with the phase variable reset.

We simulate a long orbit of this chaotic system. The system lives on a strange attractor, which can be broken into three distinct regions depending on which discontinuities points reach forwards and backwards in time. We plot in blue points which have phase reseted and are about to impact, green for points which have impacted are about to impact again and red for points which have impacted and are about to reset, see figure 7.

We apply the bump smoothing transformation to this data. The smoothed data is a single connected component. This transformation is not an injection, so for the calculation we include a delay vector in each of the original variables.

Using the state space reconstruction method presented in [5] we compute the Lyapunov spectrum of the system from the data. We calculate the largest exponent to be 0.0813, which agrees with Stefanski calculation of 0.0832. Since the system is autonomous and dissipative we know that the second largest eigenvalue is zero and that the third is negative and greater in magnitude that the first. Our results agree with this theory; we have second exponent 0.0031 and third -0.1663. Since our system analyses 6 dimensional data we could produce up to three additional spurious exponents. Our algorithm produces two further finite exponents at much larger order of

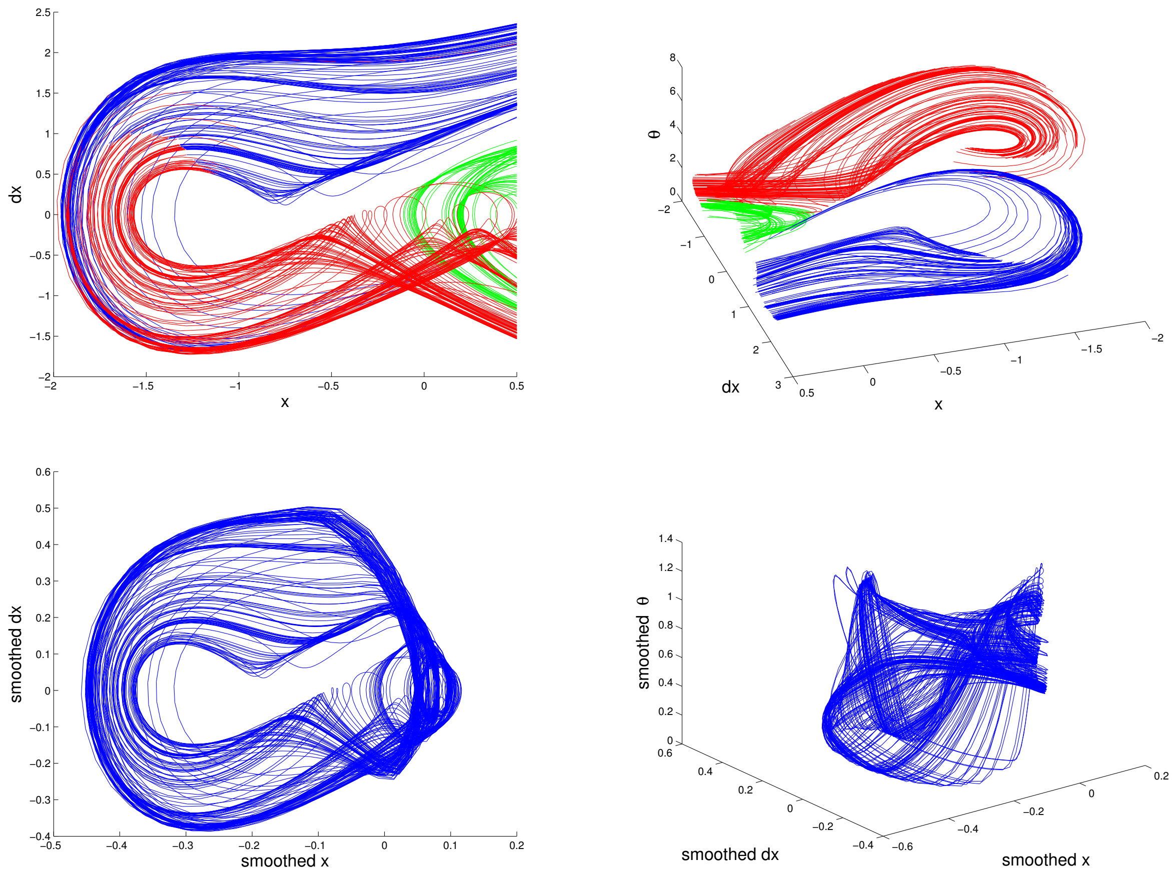

Figure 7: Top: phase and state space time series for Duffing oscillator, bottom: smoothed phase and state space time series for Duffing oscillator.

magnitude and the last exponent becomes infinite during the calculation, see figure 8.

6. Conclusion

We have shown that low-pass filters can be used to formulate smoothing transformations that map discontinuous or non-differentiable systems to smooth (or smooth everywhere except for the image of grazes) systems. We have demonstrated two different techniques for studying these smoothed systems.

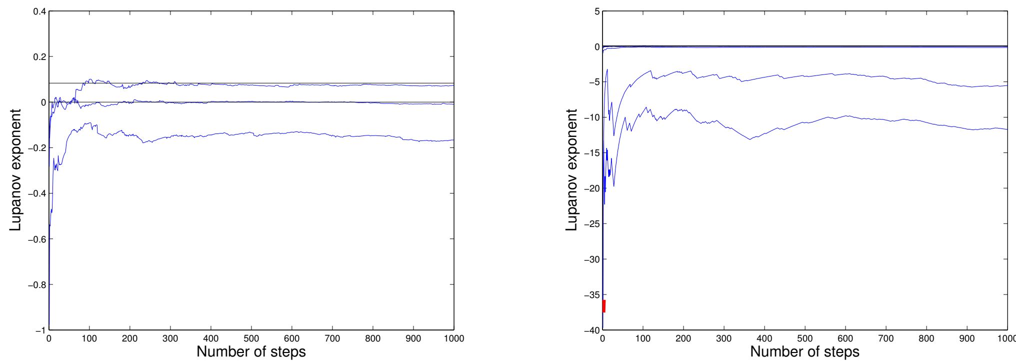

Figure 8: Left: convergence of three largest Lupanov exponents, right: behaviour of all 6 exponents, 6th exponent in bold red is taken to be $-\infty$ shortly after experiment begins.

For simple systems, we can explicitly formulate the transformation and the smoothed system to see how features in the non-smooth flow are integrated into the flow of the smoothed system.

For more complex systems we smooth a time-series, then use state space reconstruction techniques to study the smoothed system. We have shown that the smoothing procedure preserves stability properties, which gives us a novel way to calculate the Lupanov spectrum of a non-smooth system.

That these smoothing transformations are brought about whenever we apply a low-pass filter to time-series data means that experimentalists may already be inadvertently studying smooth systems of the sort presented here. Our message to experimentalists would be not to avoid this smoothing action by using filters with very short windows or by using more complex smoothing techniques such as the Savitzky Golay smoothing filter. Instead, since it is possible to understand how non-smooth features are transformed by the smoothing, they should apply low-pass filters to reduce the noise then take the effect of the smoothing into account when analysing the data. For example by looking for discontinuities in the second derivative when looking for switches after applying the moving average filter, or looking for square root behaviour in the flow when looking for grazes in a smoothed system.

In future work we hope to apply the ideas in this paper to experimental data in the spirit described above. There are many other open questions on the signal processing side of things. For instance, what is the optimal window and kernel for a smoothing filter when applied to a discontinuous system?

7. References

[1] A. Champneys P. Kowalczyk J. Guckenheimer M. d. Bernardo, C. Budd. *Piecewise-Smooth Dynamical Systems.* Springer, 2008.

[2] P. Bloomfield. *Fourier Analysis of Time Series.* Wiley, 2005.

[3] J. Huke. Embedding nonlinear dynamical systems, a guide to takens theorem. *MIMS EPrint, Manchester Institute for Mathematical Sciences, University of Manchester, UK*, 1997.

[4] J. Hook. Topics in dynamical systems (phd). *MIMS EPrint, Manchester Institute for Mathematical Sciences, University of Manchester, UK*, 2011.

[5] D. Broomhead A. Darbyshire. Robust estimation of tangent maps and liapunov spectra. *Physica D*, 89:287–305, 1996.

[6] A. Stefanski. Estimation of the largest lyapunov exponent in systems with impacts. *Chaos, Solitons, Fractals*, 15:2443–2451, 2000.
mit their bibtex database files. They are