

Signature Methods in Machine Learning

Terry Lyons & Andrew D. McLeod

January 26, 2023

Abstract

Signature-based techniques give mathematical insight into the interactions between complex streams of evolving data. These insights can be quite naturally translated into numerical approaches to understanding streamed data, and perhaps because of their mathematical precision, have proved useful in analysing streamed data in situations where the data is irregular, and not stationary, and the dimension of the data and the sample sizes are both moderate.

Understanding streamed multi-modal data is exponential: a word in n letters from an alphabet of size d can be any one of d^n messages. Signatures remove the exponential amount of noise that arises from sampling irregularity, but an exponential amount of information still remain. This survey aims to stay in the domain where that exponential scaling can be managed directly. Scalability issues are an important challenge in many problems but would require another survey article and further ideas. This survey describes a range of contexts where the data sets are small and the existence of small sets of context free and principled features can be used effectively.

The mathematical nature of the tools can make their use intimidating to non-mathematicians. The examples presented in this article are intended to bridge this communication gap and provide tractable working examples drawn from the machine learning context. Notebooks are available online for several of these examples. This survey builds on the earlier paper of Ilya Chevyrev and Andrey Kormilitzin which had broadly similar aims at an earlier point in the development of this machinery.

This article illustrates how the theoretical insights offered by signatures are simply realised in the analysis of application data in a way that is largely agnostic to the data type. Larger and more complex problems would expect to address scalability issues and draw on a wider range of data science techniques.

The article starts with a brief discussion of background material related to machine learning and signatures. This discussion fixes notation and terminology whilst simplifying the dependencies, but these background sections are not a substitute for the extensive literature they draw from.

Hopefully, by working some of the examples the reader will find access to useful and simple to deploy tools; tools that are moderately effective in analysing longitudinal data that is complex and irregular in contexts where massive machine learning is not a possibility.

Contents

1	Introduction	2
2	Background Material	4
2.1	Simple Approach to Regression	4
2.2	Transforming Datasets to Paths	6
2.3	Signature Involvement	12
2.4	Tensor Algebra and Signature	13
3	Signature as a Feature Map	15
4	Controlled Differential Equations and the Log-ODE Method	18
5	Computing Signatures	20
6	Expected Signature	23

7	Truncation Order Selection	25
8	Signature Kernel	27
9	Neural CDEs and Neural RDEs	31
10	Speech Emotion Recognition	34
11	Health Applications	36
11.1	Bipolar and Borderline Personality Disorders	36
11.2	Alzheimer’s Disease	41
11.3	Early Sepsis Detection	42
11.4	Information Extraction from Medical Prescriptions	45
12	Landmark-based Human Action Recognition	49
12.1	Path Disintegration and Transformations	50
12.2	Datasets	51
12.3	Method Implementation	51
12.4	Comparison with Specifically-Tailored Methods	54
12.5	Demo Notebook	55
13	Distribution Regression via the Expected Signature	55
14	Anomaly Detection	58

1. Introduction

Streams of sequential information can be found almost everywhere in our society. They need to be understood, analysed, and interpreted. It might be a stream of spoken words, financial data, or ones hospital records. The simplest example of sequential information is probably a text. One sees quickly that the number of words is exponential in the number of letters. Most streams involve multiple channels, and understanding the consequences of the “data” involve understanding how the channels interleave. Suppose one channel recorded arrivals of stock from deliveries at a supermarket. Another involved the demand from customers. To predict sales requires a detailed understanding of how these channels interact. The order of events matters. This can be very expensive using classical methodology. Smoothing the deliveries and the demand result in data that does not indicate the shortfall because the demand came before the arrival of deliveries.

The signature of a path is a forgetful summary that is intrinsically multidimensional. It captures the order in which things happen, the order in which the path visits locations, but completely ignores the parameterisation of the data. It does not matter if the path was sampled 75 times because there were 75 deliveries that week and 65 times the week after. But it remembers the order of events perfectly. In a natural way, it produces a vector embedding of the path that is not sensitive to how often the path is sampled but is fully sensitive to the order of events on the different channels. This can allow the detection of such events fare more efficiently than the customary time series approaches. The signature of a path is a mathematical transform that faithfully describes the curve, or the unparameterised path, but removes completely the infinite dimensional group of parameterisation based data symmetries.

In this article, we survey a variety of contexts where signatures have been applied to machine learning. This survey is intended to act as a pseudo-sequel to the introduction provided in [CK16]. The applications selected for inclusion illustrate that the theoretical benefits, guaranteed by the underlying theory of path signatures, are being fully realised in practical problems. In some sense, our article may be thought of as a continuation of Section 2 in [CK16] to cover more sophisticated uses of the signature transform in machine learning made after the appearance of [CK16].

In the interest of readability we do not present this article as a proper sequel to [CK16]. Instead, we cover some over-lapping material within our first few sections, hence the term pseudo-sequel. The inclusion of such material is two-fold. Firstly, they provide a minimal coverage of the material required for the later surveys of the applications of signature techniques for the reader unfamiliar with the basic ideas, making this work more

self-contained. Secondly, they allow us to fix notation and terminology that will be used later, forming a sort-of dictionary that can be referred back to during the latter sections.

The over-lapping material is approached from a different perspective to the presentation in [CK16]. Section 2.1 introduces regression from the viewpoint of a beginner who has no prior experience of machine learning. A few basic strategies are covered in Sections 2.1 and 2.2 that naturally lead to the signature transforms informal introduction in Section 2.3. The machine learning material included in Section 2.1 is only the essentials we later require. The reader seeking a more comprehensive introduction to machine learning is directed to the book [Gér19]. Similarly, the transformations of datasets to paths presented in Section 2.2 only cover those used in later applications. A more thorough introduction to such transformations and their impact on signatures can be found in [CK16].

After informally introducing the signature as a method to summarise a path in Section 2.3 that captures all the relevant information, the mathematical framework required to rigorously talk about signatures is covered in Section 2.4. Keeping in mind our adopted machine learning perspective, signature properties and related theorems are all stated for continuous paths of bounded variation. Indeed we stick to the convention of only considering continuous paths of bounded variation throughout the article, making only the occasional brief mention (without justification) of a properties extension to less regular *rough paths*.

As a pseudo-sequel to [CK16], we assume the reader has had some exposure to signatures. Our discussion of signatures assumes a greater mathematical maturity than [CK16] and, in contrast to the example-heavy introduction given in [CK16], we favour a succinct presentation of the key theoretical aspects. Only the mathematical properties of signatures directly relevant to the later applications are included in Section 2.4. A wider range of properties of path signatures can be found in the survey article [Lyo14], whilst the reader seeking a comprehensive introduction to the theory is directed to the lecture notes [LCL04].

The reader already comfortable with *Regression, Feature Maps, Kernels, Embeddings of datasets into paths*, and the mathematical theory of *Signatures* may safely skim Sections 2.1 - 2.4, using them only as a dictionary for notation and terminology as required when reading the later sections. The material covered ceases over-lapping with the content of [CK16] in Section 3. For the readers ease of navigation, we briefly summarise the content included in the remaining sections.

- Section 3 provides the theoretical results underpinning the suitability of the signature transform as a feature map.
- Section 4 introduces *Controlled Differential Equations* (CDEs) and covers the *log-ODE* method. The log-ODE method is a technique for numerically solving CDEs using the *log signature*, which is itself introduced in Section 2.4.
- Section 5 provides an introduction to the *Esig* python package [Lyo02] for computing truncated signatures and log signatures.
- Section 6 explains how signatures can be extended to distributions *without* any restrictions on the distributions, as in [CO18] for example. We additionally outline how this theory offers an approach to distribution regression following [BDLLS20].
- Section 7 covers the work [ABKLS19] proposing how the signature transform may be incorporated into a neural network. A particular consequence is the proposal that the signature truncation level should be learnt, and hence be determined by the data.
- Section 8 focuses on the fact that the signature transform is a kernel. It covers the work of [CFLSY20] detailing how the associated kernel function can be effectively approximated. This allows the full signature transform to be used *without* truncation.
- Section 9 covers the introductions of *Neural Controlled Differential Equations* (Neural CDEs) made in [FKLM20] and of *Neural Rough Differential Equations* (Neural RDEs) made in [FKLMS21]. The universality guarantees associated with Neural CDEs are fundamentally reliant on the theory of path signatures, whilst Neural RDEs are built around the theory of log signatures.
- Section 10 covers the work of [LLNNSW19] on how signature transform methods can be used in the extraction of emotion from segments of speech.

- Section 11 covers a range of medical applications of signature transform methods. These include distinguishing between Bipolar Disorder and Borderline Personality Disorder [AGGLS18, LLNSTWW20], diagnosing Alzheimer’s disease [GLM19], information extraction from medical prescriptions [BLNW20], and early detection of Sepsis [HKLMNS19, HKLMNS20].
- Section 12 covers the work of [JLNSY17] using signature transform methods to recognise human actions from the motion of 20-40 markers located on the person. A python notebook exists (see Subsection 12.5) allowing the reader to train their own classifier (following a variant of the method proposed in [JLNSY17]) within 20 minutes on standard hardware (including downloading relevant files and carrying out training).
- Section 13 covers the work of [BDLLS20] realising the use of a generalised signature transform for distribution regression both with and without truncation of the involved signatures.
- Section 14 covers the use of signatures for anomaly detection proposed in [ACFL20].

Acknowledgements: Terry Lyons was supported in part by the DataSig programme funded by the EPSRC [grant number EP/S026347/1], in part by The Alan Turing Institute under the EPSRC grant EP/N510129/1, the Data Centric Engineering Programme (under the Lloyd’s Register Foundation grant G0095), the Defence and Security Programme (funded by the UK Government) and the Office for National Statistics & The Alan Turing Institute (strategic partnership) and in part by the Hong Kong Innovation and Technology Commission (InnoHK Project CIMDA). Andrew McLeod was supported by the DataSig programme funded by the EPSRC [grant number EP/S026347/1].

2. Background Material

2.1. Simple Approach to Regression

Regression is a basic, yet extremely common, learning task in which real-valued functions are modelled using sets of data. The fundamental goal of regression is to learn a function that predicts outcome values based on a specified collection of input features. We discuss the mathematical framework associated to simple regression tasks. For examples of working regression into machine learning pipelines see [Gér19].

We adopt the convention that a *dataset* will refer to a finite subset of a Banach space, whilst a *labelled dataset* will refer to a finite subset of $V \times \mathbb{R}$ where V is a Banach space. If $\Omega := \{(x_i, y_i) : i \in \mathcal{I}\} \subset V \times \mathbb{R}$ is a labelled dataset then for each $i \in \mathcal{I}$ the value $y_i \in \mathbb{R}$ represents the systems response to the inputs $x_i \in V$. We denote by Ω_V the dataset given by the projection of Ω to V , that is $\Omega_V := \{x_i : i \in \mathcal{I}\}$. The task is to transition from the collection of pairs Ω to a function $f \in C^0(V; \mathbb{R})$ that can be used to predict the systems response to input features $x \in V$ that do *not* correspond to an instance within Ω_V .

In this formulation, we are implicitly assuming that the dataset Ω_V gives a “good approximation” (in some sense) of the entirety of V . That is, we assume that it is reasonable to try and learn the systems response to inputs in V using only its responses to inputs in Ω_V . This restriction is natural within learning tasks.

We would not expect a house value predictor learnt from rural house data to work well on houses in all locations. Nor would we expect a reaction time predictor learnt from professional racing drivers to do a good job at predicting the reaction time of members of the general public. And it is unlikely that knowing the reaction of mice to a drug would allow the accurate prediction of the reaction of all animals to the drug.

The issue with all these examples is the presence of elements in V bearing little resemblance to *any* element in Ω_V . There are inner city houses with very little similarity to rural houses, there are members of the general population sharing few similar characteristics with professional racing drivers, and there are animals whose biological make up is significantly different to that of mice. It is not reasonable to expect to accurately predict a systems response to inputs bearing little resemblance to the inputs for which the response is known.

It is easy to imagine several *ad-hoc* methods of ensuring that every element in V is somehow similar to an element in Ω_V . But introducing a mathematical framework capable of describing similarity is more challenging. We will later discuss a rigorous notion of similarity (cf. Section 14). But for now, we just assume

that Ω_V provides a "sufficiently good" approximation of V , without concerning ourselves with the precise details meant by "sufficiently good".

A core idea in many schemes for learning a function f from a collection of pairs Ω is to identify basic functions or features that are readily evaluated at each point in Ω_V . Once identified, we try to express the observed function as a linear combination of these basic functions.

To illustrate this idea, suppose we want to learn a continuous function $\rho : [0, 1] \rightarrow \mathbb{R}$ using its values at a given collection of points $x_1, \dots, x_N \in [0, 1]$, for some $N \in \mathbb{N}$. To match the terminology introduced above, we have a labelled dataset $\{(x_i, y_i := \rho(x_i)) : i \in \{1, \dots, N\}\} \subset [0, 1] \times \mathbb{R}$, we assume that the points x_1, \dots, x_N give a "good approximation" (in some sense) of $[0, 1]$, and that we want to learn the response of ρ to any input in $[0, 1]$.

An elementary approach is to consider the functions $\phi_k : [0, 1] \rightarrow \mathbb{R}$ for $k \in \mathbb{N}$, defined by $\phi_k(x) := x^k$, as the feature functions, i.e. the monomials. Then we fix a (large) $K_0 \in \mathbb{N}$ and try to express

$$\rho \approx \sum_{k=0}^{K_0} a_k \phi_k. \quad (2.1)$$

If (2.1) is possible, then the coefficients a_k are given by the linear equations

$$y_i = \sum_{k=0}^{K_0} a_k \phi_k(x_i) \quad (2.2)$$

for each $i \in \{1, \dots, N\}$. We want to use the linear combination in (2.1) to predict the value of ρ at new instance in $[0, 1]$. If the equations (2.2) are non-degenerate, then the solution is unstable; any new instance will lead to a completely different set of coefficients a_k . Hence it is desirable that these equations should be degenerate, and solving (2.2) often involves numerical techniques such as *Singular Value Decomposition* (SVD) or *gradient descent* algorithms involving suitable cost functions.

The success of this approach to approximating ρ is fundamentally reliant on the linear span of the set $\{\phi_k : k \in \mathbb{N}\}$ being dense within the class of functions $C^0([0, 1]; \mathbb{R})$. In this case, since the monomials span an algebra and the unit interval is a compact subset of \mathbb{R} , the *Stone-Weierstrass theorem* (see [Sto48], for example) provides the required density.

Returning to a general dataset $\Omega_V \subset V$, one approach to finding a suitable collection of feature functions is to identify a *universal feature map* Ψ . That is, to find an embedding $\Psi : V \rightarrow W$ for a Banach space W for which

$$\overline{\Psi(V)^*} = C^0(\Psi(V)). \quad (2.3)$$

In words, a universal feature map Ψ provides a representation of the dataset Ω_V as a subset $\Omega_W := \Psi(\Omega_V)$ of a Banach space W in which continuous functions can be approximated by linear functions. This would allow us to use linear functions as our feature functions. Before mimicking the approach of approximating a continuous function on the interval $[0, 1]$, we need to understand how to evaluate $\phi(x_i)$ for elements $\phi \in \Psi(X)^*$ and $i \in \mathcal{I}$.

We seek a suitable universal embedding of the set V into some Banach space W . If we simultaneously found an embedding of V into the dual space W^* , then we would have what is called a *kernel* in machine learning. To be more precise, a *Kernel* is a pair of embeddings

$$\Psi : V \rightarrow W \quad \text{and} \quad \Phi : V \rightarrow W^*. \quad (2.4)$$

A kernel allows one to define a function $K : V \times V \rightarrow \mathbb{R}$ by setting

$$K(x, y) := \langle \Psi(x), \Phi(y) \rangle \quad (2.5)$$

where $\langle \cdot, \cdot \rangle : W \times W^* \rightarrow \mathbb{R}$ is the natural dual-pairing map. That is, given $v \in W$ and $\phi \in W^*$ we define $\langle v, \phi \rangle := \phi(v)$. Sometimes the function K defined in (2.5) is itself referred to as a kernel. In the literature it is common to require W to be a Hilbert space, in which case the dual space W^* is isometric to W itself, and the natural dual-pairing map is simply the inner product on W .

One major benefit of kernels is that they provide a concrete approach to the learning problem. We can use

the kernel functions $K(\cdot, x_i) = \langle \cdot, \Phi(x_i) \rangle \in V^*$ as the feature functions. That is, returning to the observed labelled dataset Ω , first solve the system

$$y_i = \sum_{k \in \mathcal{I}} a_k K(x_i, x_k) \quad (2.6)$$

for each $i \in \mathcal{I}$, and then express the observed function f as

$$f(\cdot) := \sum_{k \in \mathcal{I}} a_k K(\cdot, x_k). \quad (2.7)$$

Neither the dimension of W or the explicit form of the embeddings Ψ and Φ are involved. The strategy only involves a *finite* number of terms, namely the $|\mathcal{I}|^2$ numbers $K(x_i, x_j)$, which is independent of the dimension of W . Assuming K is known, computing these $|\mathcal{I}|^2$ numbers does not require computing the embedded data points $\Phi(x_i)$ and $\Psi(x_i)$ themselves. Consequently there is no restriction on the embeddings that may be considered, with embeddings into infinite dimensional vector spaces W being allowed both theoretically and in practice. But the number of computations does grow quadratically with respect to the size of the dataset Ω_V , which becomes problematic for large datasets Ω_V .

Another limitation is that a kernel K does not necessarily provide a *graded* set of feature functions. A priori, there is no natural notion of ordering on the set $\{K(\cdot, x_i) : i \in \mathcal{I}\}$. Returning to the monomials on $[0, 1]$, namely the set $\{\phi_k : [0, 1] \rightarrow \mathbb{R}, \phi_k(x) := x^k : k \in \mathbb{N}_0\}$, then there is a notion of grading. This is given by the order of the monomial, with there being an intrinsic idea that using lower order monomials to approximate a function ρ is somehow better (or simpler) than using higher order monomials. For example, it is preferable to use ϕ_2 in place of ϕ_{10} if possible. The feature functions provided by a generic kernel K are not guaranteed to have any such notion. Given distinct $i, j \in \mathcal{I}$, it is not clear whether there is any benefit between using $K(\cdot, x_i)$ or $K(\cdot, x_j)$ to approximate a function.

2.2. Transforming Datasets to Paths


A finite dataset $\Omega_V \subset V$ is a *stream*, that is, an element of

$$\mathcal{S}(V) := \{(x_1, \dots, x_k) : k \in \mathbb{N}, x_j \in V\}. \quad (2.8)$$

The length of a stream $\Sigma \in \mathcal{S}(V)$ is given by its cardinality $\#(\Sigma)$. For each point $x_i \in \Omega_V$ we know that the systems response is given by $y_i \in \mathbb{R}$. We want to summarise the stream to succinctly capture its effects. This can be achieved by transforming the dataset Ω_V to a *path* (a continuous map $\gamma : I \rightarrow V$ where $I \subset \mathbb{R}$ is an interval) and summarising the resulting path.

There are numerous different ways in which a stream can be transformed into a path and several explicit examples of such transformations may be found in, for example, [CK16]. For the readers convenience, we briefly summarise some of the more important examples for our purposes. A relatively simple class of transformations involve choosing an order in which the points of Ω_V are visited and specifying how two successive points are connected. If the dataset Ω_V consists of sequential data points then there is a natural pre-determined ordering; if not, then a choice of ordering must be made.

Two of the simplest transformations are given by linear interpolation (where each point is joined to the next via a straight line) and rectilinear interpolation (where each point is joined to the next via coordinate axis lines). These transformations are illustrated below in Figure 1 in the setting that $V := \mathbb{R}^2$ and we consider the stream $\Omega_e := \{(1, 0), (4, 1), (2, 3), (3, 5), (5, 2), (6, 7)\} \in \mathcal{S}(\mathbb{R}^2)$.

Figure 1: Examples of different transformations of Ω_e to a path in \mathbb{R}^2

There are numerous alternative ways to interpolate between points that could be considered, each giving rise to different transformations.

A more involved class of transformations involve *augmenting* the stream Ω_V before applying a choice of interpolation between the resulting points to provide a path. More precisely, we consider a fixed map $\Theta : V \rightarrow W$ for some Banach space W , and interpolate the augmented dataset $\Theta(\Omega_V) \subset W$. We now introduce some of the more commonly used explicit augmentations. The *Lead-Lag* augmentation is the map $T_{\text{Lead-Lag}} : \mathcal{S}(V) \rightarrow \mathcal{S}(V \times V)$ defined by

$$T_{\text{Lead-Lag}}(v_1, \dots, v_k) := ((v_1, v_1), (v_2, v_1), (v_2, v_2), \dots, (v_j, v_j), (v_{j+1}, v_j), \dots, (v_k, v_k)). \quad (2.9)$$

It takes a stream of length k to a stream of length $2k$, whilst splitting the original stream into two copies, labelled the future and the past. There is a delay between when the future is updated and when the past is subsequently updated. The length of delay can be varied, leading to a Lead-Lag where the number of steps between the future being updated and the past being updated is at least two. Further, more than one past stream can be recorded, with each past stream having a different number of delay steps between when the future stream is updated and when it is itself updated.


Take $V := \mathbb{R}$ and consider the stream $\Omega_L := \{1, 6, 3, 9, 5\}$. The basic Lead-Lag augmentation defined in (2.9) would map this stream to the stream

$$\{(1, 1), (6, 1), (6, 6), (3, 6), (3, 3), (9, 3), (9, 9), (5, 9), (5, 5)\} \in \mathcal{S}(\mathbb{R}^2).$$


If the delay is increased to 2 steps, then we get the stream

$$\{(1, 1), (6, 1), (3, 1), (3, 6), (9, 6), (5, 6), (5, 3), (0, 3), (0, 3), (0, 9), (0, 9), (0, 9), (0, 5)\} \in \mathcal{S}(\mathbb{R}^2).$$

These streams can then be interpolated to give paths in \mathbb{R}^2 as we illustrate below in Figure 2.



(a) Piecewise Linear Interpolation of Lead-Lag augmentation with delay of one step



(b) Piecewise Linear Interpolation of Lead-Lag augmentation with delay of two steps

Figure 2: Examples of Lead-Lag transformations with differing lengths of delay

We could record both the one-step delayed past and the two-step delayed path in a single augmentation. This would map Ω_L to the stream

$$\left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 6 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 6 \\ 6 \\ 1 \end{pmatrix}, \begin{pmatrix} 6 \\ 6 \\ 6 \end{pmatrix}, \begin{pmatrix} 3 \\ 6 \\ 6 \end{pmatrix}, \begin{pmatrix} 3 \\ 3 \\ 6 \end{pmatrix}, \begin{pmatrix} 3 \\ 3 \\ 3 \end{pmatrix}, \begin{pmatrix} 9 \\ 3 \\ 3 \end{pmatrix}, \begin{pmatrix} 9 \\ 9 \\ 3 \end{pmatrix}, \begin{pmatrix} 9 \\ 9 \\ 9 \end{pmatrix}, \begin{pmatrix} 5 \\ 9 \\ 9 \end{pmatrix}, \begin{pmatrix} 5 \\ 5 \\ 9 \end{pmatrix}, \begin{pmatrix} 5 \\ 5 \\ 5 \end{pmatrix} \right\} \in \mathcal{S}(\mathbb{R}^3).$$

This stream can be interpolated to give a path in \mathbb{R}^3 as illustrated below in Figure 3.




Figure 3: Multi-past Lead-Lag Transformation of Ω_L

The final variant of this transformation we mention involves not pausing the future stream whilst the past streams are updated. If d past streams are being recorded, then the *no future pause* augmentation maps a

stream $(v_1, \dots, v_k) \in \mathcal{S}(V)$ to the stream

$$\left\{ \begin{pmatrix} v_1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} v_2 \\ v_1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \dots, \begin{pmatrix} v_{d+1} \\ v_d \\ v_{d-1} \\ \vdots \\ v_2 \\ v_1 \end{pmatrix}, \begin{pmatrix} v_{d+2} \\ v_{d+1} \\ v_d \\ \vdots \\ v_3 \\ v_2 \end{pmatrix}, \dots, \begin{pmatrix} v_k \\ v_{k-1} \\ v_{k-2} \\ \vdots \\ v_{k-d+1} \\ v_{k-d} \end{pmatrix}, \begin{pmatrix} 0 \\ v_k \\ v_{k-1} \\ \vdots \\ v_{k-d+2} \\ v_{k-d+1} \end{pmatrix}, \dots, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ v_k \\ v_{k-1} \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ v_k \end{pmatrix} \right\} \quad (2.10)$$

in $\mathcal{S}(V^{d+1})$. Returning to the dataset Ω_L , applying the no future pause augmentation defined in (2.10) with $d = 2$ results in the new stream

$$\left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 6 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 6 \\ 1 \end{pmatrix}, \begin{pmatrix} 9 \\ 3 \\ 6 \end{pmatrix}, \begin{pmatrix} 5 \\ 9 \\ 3 \end{pmatrix}, \begin{pmatrix} 0 \\ 5 \\ 9 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 5 \end{pmatrix} \right\} \in \mathcal{S}(\mathbb{R}^3).$$

The piecewise linear interpolation of this stream results in the path in \mathbb{R}^3 illustrated in Figure 4.




Figure 4: No Future Pause Transformation of Ω_L

Another common augmentation is the *Time* augmentation $T_{\text{Time}} : \mathcal{S}(V) \rightarrow \mathcal{S}(V \times \mathbb{R})$ defined by

$$T_{\text{Time}}(v_1, \dots, v_k) := ((v_1, t_0), \dots, (v_j, t_j), \dots, (v_k, t_k)) \quad (2.11)$$

for a strictly increasing sequence of times $0 \leq t_1 < t_2 < \dots < t_k$, ensuring the resulting path has a strictly monotonic component. A variant of this is the *Time-difference* augmentation $T_{\text{Time-diff}} : \mathcal{S}(V) \rightarrow \mathcal{S}(V \times \mathbb{R})$ defined by

$$T_{\text{Time-diff}}(v_1, \dots, v_k) := ((v_1, t_0), \dots, (v_j, t_j - t_{j-1}), \dots, (v_k, t_k - t_{k-1})) \quad (2.12)$$

for a strictly increasing sequence of times $0 \leq t_1 < t_2 < \dots < t_k$. Both the Time and Time-diff augmentations take streams of length k to streams of length k .

To illustrate these augmentations, consider the stream $\Omega_T := \{1, 5, 2, 9, 7, 6\} \in \mathcal{S}(\mathbb{R})$ and the strictly


increasing sequence 0, 1, 3, 6, 8, 12. Then

$$T_{\text{Time}}(\Omega_T) = \{(1, 0), (5, 1), (2, 3), (9, 6), (7, 8), (6, 12)\} \in \mathcal{S}(\mathbb{R}^2) \quad (2.13)$$


and

$$T_{\text{Time-diff}}(\Omega_T) = \{(1, 0), (5, 1), (2, 2), (9, 3), (7, 2), (6, 4)\} \in \mathcal{S}(\mathbb{R}^2). \quad (2.14)$$

In Figure 5 we illustrate the paths resulting from both the piecewise linear and rectilinear interpolations of these streams.



(a) Time augmentation combined with both piecewise linear interpolation (blue line) and rectilinear interpolation (yellow line)



(b) Time-diff augmentation combined with both piecewise linear interpolation (blue line) and rectilinear interpolation (yellow line)

Figure 5: Example Time and Time-diff transformations of Ω_T to a path in \mathbb{R}^2

The final explicit augmentation we mention is the *Invisibility Reset* augmentation $T_{\text{inv}} : \mathcal{S}(V) \rightarrow \mathcal{S}(V \times \mathbb{R})$ defined by

$$T_{\text{inv}}(v_1, \dots, v_k) := ((v_1, 1), \dots, (v_j, 1), \dots, (v_k, 1), (v_k, 0), (0, 0)) \quad (2.15)$$

and takes streams of length k to streams of length $k + 2$. The resulting stream now contains information on the norm of the original stream elements v_1, \dots, v_k in a translation-invariant manner. We apply this augmentation to the one-dimensional stream $\Omega_{ir1} := \{1, 3, 4, 8, 9\}$ in $\mathcal{S}(\mathbb{R})$ and the two-dimensional stream $\Omega_{ir2} := \{(1, 2), (3, 4), (4, 6), (5, 9), (7, 10)\}$ in $\mathcal{S}(\mathbb{R}^2)$. This yields

$$T_{\text{inv}}(\Omega_{ir1}) = \{(1, 1), (3, 1), (4, 1), (8, 1), (9, 1), (9, 0), (0, 0)\} \in \mathcal{S}(\mathbb{R}^2) \quad (2.16)$$

and

$$T_{\text{inv}}(\Omega_{ir2}) = \left\{ \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ 4 \\ 1 \end{pmatrix}, \begin{pmatrix} 4 \\ 6 \\ 1 \end{pmatrix}, \begin{pmatrix} 5 \\ 9 \\ 1 \end{pmatrix}, \begin{pmatrix} 7 \\ 10 \\ 1 \end{pmatrix}, \begin{pmatrix} 7 \\ 10 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \right\} \in \mathcal{S}(\mathbb{R}^3). \quad (2.17)$$

Examples of paths resulting from interpolating these new streams are illustrated below in Figures 6 and 7.




Figure 6: Invisibility Reset Augmentation combined with piecewise linear interpolation applied to Ω_{ir1}




Figure 7: Invisibility Reset Augmentation combined with both piecewise linear interpolation (blue line) and rectilinear interpolation (black line) applied to Ω_{ir2}

2.3. Signature Involvement

Given the numerous different possible ways of transforming a stream to a path, see Section 2.2 for some explicit examples, there is evidently no universal choice that will always be the best option. The choice of transformation is a modelling problem; it requires making an assumption about how the system should change on the arrival of the next instance. Lead-Lag type transformations are frequently used in financial markets modelling since they can encode the modelling assumption that no strategy should be able to use the present time stock value. Strategies are only allowed to use the past streams of the augmented dataset. Since the future stream is always updated first, with there being a definite delay before any of the past streams are subsequently updated, the strategies are prevented from making use of the present time stock values.

Once we have appropriately transformed a dataset Ω_V to a path $\gamma : I \rightarrow V$, we seek to summarise it (throw away irrelevant information) so as to capture its effect. It is instructive to realise that the classical mathematical description of paths is ill-suited to providing such a summary. The classical description requires precisely recording the position of the path at every instance. Even the probabilistic Kolmogorov approach (in which one considers fixed times t_i , open sets O_i , and considers the probability that, for all i , at time t_i the path is within the set O_i) still places emphasis on the location at a specified time. The limitations in using this approach to summarise paths become clear from considering some examples. We would not try to summarise a movie by recording a second-by-second account, we would not describe a football match by providing the exact position of the ball at each second and we would not describe a drawing by providing a second-by-second account of the location of the pencil tip. The problem is that we are not throwing away any irrelevant information, making the extraction of the useful information very challenging.

Recording only the major events of a path rather than all the events may seem like a solution that could be achieved via sampling; however this still run into problems. The issue is that the order in which the major events occur is frequently at least as important as the major events themselves. Suppose both the pulse and breathing rate of a patient are recorded and both are discovered to have decreased. From a diagnostic perspective it could be critical to know which decreased first. The gap between the two events may be very small, meaning that a very high frequency sampling rate of both the pulse and the breathing rate is required to capture their order. But without a priori knowledge of when the first drop occurs we have to sample the entirety of both streams, resulting in a large amount of irrelevant information being recorded. This problem is only amplified as the number of events $n \geq 2$ whose order we are interested in increases. We want a summary of the main events and the order in which they happen that is detailed enough to distinguish between different paths whilst still discarding irrelevant information.

We would also like our summary to be invariant under reparameterisation of the path. The response of a system to a given path will not be affected by the speed at which we are moving along the path; the features of the path will induce the same response regardless of how quickly we choose to experience them. For example, consider writing the character ‘3’. The effect of this path is a figure ‘3’ drawn on our piece of paper, and obviously this effect remains the same irrespective of how quickly we draw it. Two 3’s drawn with different speeds are still both a ‘3’; the change of speed with which we traverse the path given by ‘writing a figure 3’ does *not* change its effect in producing a ‘3’ on our page.

The freedom of reparameterisation makes the embedding of a stream into a continuous path a reasonable idea. Each point in a stream represents a change in the state of the system and we can certainly choose an order in which they happen (determined by the modelling assumptions). This results in a system whose state evolves over time, i.e. a path. The concern is that these changes appear to happen instantaneously, seemingly giving a discontinuous path. But via reparameterisation, we can imagine slowing the speed down and introducing new virtual time over which the next change to the system happens. By sufficiently slowing the speed, we can imagine that the seemingly instantaneous, with respect to the original time, affect of the next instance becomes a continuous change with respect to the new virtual time.

Recalling our machine learning motivations, after embedding a dataset into a continuous path, our problem has been transformed to seeking to learn a function that is an effect of this path. A consequence of *rough path theory* is that the *signature* provides a summary of the paths affect on systems that satisfies our list of requirements. The signature of a path determines this path in an essentially unique way and is invariant under reparameterisation of the path. This invariance allows the signature to readily record the order in which events occur without recording precisely when they occur. The signature of a path even provides a natural feature set of linear functionals capturing the aspects of the data necessary to predict the effects of the path on systems.

The signature of a path x is the response of the exponential nonlinear system to the path. In one dimension the response of the exponential linear system captures the monomials, in the sense that the power series of $\exp : \mathbb{R} \rightarrow \mathbb{R}$ is given by

$$\exp(x) = \sum_{k=0}^{\infty} \frac{x^k}{k!}. \quad (2.18)$$

The projection onto the m^{th} degree term gives the m^{th} degree monomial. Thus, at least heuristically, we see why we expect the signature to be a suitable feature map. In Section 3, we provide some of the details justifying these assertions. But first we must introduce the space upon which the signature will live.

2.4. Tensor Algebra and Signature

Let V be a d -dimensional Banach space with basis $\mathcal{B} = \{v_1, \dots, v_d\}$ and denote the dual basis of V^* by $\mathcal{B}^* = \{v_1^*, \dots, v_d^*\}$. The Tensor algebra of V is given by

$$T(V) := \bigoplus_{k=0}^{\infty} V^{\otimes k} \quad (2.19)$$

where $V^{\otimes 0} := \mathbb{R}$. Denote the projection $T(V) \rightarrow V^{\otimes n}$ by π_n for each $n \in \mathbb{N}_0$. We can consider the truncated tensor algebra

$$T^{(n)}(V) := \bigoplus_{k=0}^n V^{\otimes k} \quad (2.20)$$

and, for each $n \in \mathbb{N}_0$, we denote the projection $T(V) \rightarrow T^{(n)}(V)$ by Π_n . The tensor product \otimes extends to $T(V)$ by defining

$$A \otimes B := \bigoplus_{n=0}^{\infty} \sum_{k=0}^n A_k \otimes B_{n-k} \quad (2.21)$$

when $A, B \in T(V)$ are given by $A = \bigoplus_{n=0}^{\infty} A_n$ and $B = \bigoplus_{n=0}^{\infty} B_n$.

For each $n \in \mathbb{N}$ the basis \mathcal{B} determines the basis

$$\mathcal{B}^{\otimes n} := \{v_{\mathbf{K}} = v_{k_1} \otimes \dots \otimes v_{k_n} : \mathbf{K} = (k_1, \dots, k_n) \in \{1, \dots, d\}^n\} \quad (2.22)$$

for $V^{\otimes n}$ and the corresponding dual basis

$$(\mathcal{B}^*)^{\otimes n} := \{v_{\mathbf{K}}^* = v_{k_1}^* \otimes \dots \otimes v_{k_n}^* : \mathbf{K} = (k_1, \dots, k_n) \in \{1, \dots, d\}^n\} \quad (2.23)$$

for $(V^*)^{\otimes n}$.

For $n \in \mathbb{N}$ and $\phi \in (V^*)^{\otimes n} = (V^{\otimes n})^*$, we can lift ϕ to an element of $T(V)^*$ by defining $\phi(A) := \phi(\pi_n(A))$ for $A \in T(V)$. Thus $(V^*)^{\otimes n} \subset T(V)^*$ for each $n \in \mathbb{N}$, and so we have $T(V^*) \subset T(V)^*$ by linearity. In particular, if $\phi = \bigoplus_{n=0}^{\infty} \phi_n \in T(V^*)$ and $v = \bigoplus_{n=0}^{\infty} v_n \in T(V)$, then $\phi(v) = \sum_{n=0}^{\infty} \phi_n(v_n)$.

The choice of basis \mathcal{B} induces an inner product on V by defining

$$\langle v_i, v_j \rangle_V := \delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \quad (2.24)$$

and extending bilinearly to the whole of V . Given any $n \in \mathbb{N}$ the inner product defined in (2.24) can be extended to $V^{\otimes n}$ by defining

$$\langle v_{i_1} \otimes \dots \otimes v_{i_n}, v_{j_1} \otimes \dots \otimes v_{j_n} \rangle_{V^{\otimes n}} := \prod_{k=1}^n \langle v_{i_k}, v_{j_k} \rangle_V = \delta_{i_1 j_1} \dots \delta_{i_n j_n} \quad (2.25)$$

and again extending bilinearly to the whole of $V^{\otimes n}$. This inner product is naturally extended to $T(V)$ by

setting

$$\langle A, B \rangle_{T(V)} = \sum_{n=0}^{\infty} \langle \pi_n(A), \pi_n(B) \rangle \quad (2.26)$$

for $A, B \in T(V)$. Thus $(T(V), \langle \cdot, \cdot \rangle)$ is an inner product space, hence the completion of $T(V)$, denoted by $\overline{T(V)}$, is a Hilbert space when equipped with this inner product. In terms of this inner product, for every $n \in \mathbb{N}$, the elements of the dual basis (2.23) are given by $v_K^*(\cdot) = \langle \cdot, v_K \rangle$.

Now suppose that x is a stream of information defined on $[a, b]$ with features in V . Then the *signature* of x is the solution to the universal differential equation

$$dS_{a,t}(x) = S_{a,t}(x) \otimes dx_t \quad \text{with} \quad S_{a,a}(x) = 1 = 1 \oplus 0 \oplus 0 \oplus \dots \in T(V), \quad (2.27)$$

formalising that the signature $S(x)$ is the response of the exponential nonlinear system to the stream x .

In the case that $x : [a, b] \rightarrow V$ is a continuous path of bounded variation, then the signature of x is given by the iterated integrals of x . That is, $S_{a,t}(x) = \bigoplus_{n=0}^{\infty} S_{a,t}^n(x)$ where

$$S_{a,t}^0(x) \equiv 1 \quad \text{and} \quad S_{a,t}^n(x) = \int_{a < t_1 < \dots < t_n < t} \dots \int dx_{t_1} \otimes \dots \otimes dx_{t_n} \quad \text{for} \quad n \geq 1. \quad (2.28)$$

One can still make sense of the signature of a less regular path x , see [LCL04], for example.

We now state some properties of signatures that are particularly useful from our machine learning perspective. The signature $S(x)$ is a *multiplicative* functional $\Delta_{[a,b]} \rightarrow T(V)$, where $\Delta_{[a,b]} := \{(s, t) : a \leq s \leq t \leq b\}$. By multiplicative, we mean that Chen's identity holds; namely,

$$S_{s,u}(x) \otimes S_{u,t}(x) = S_{s,t}(x) \quad (2.29)$$

whenever $a \leq s \leq u \leq t \leq b$. A consequence of (2.29) is that the signature of a concatenated path decomposes into the tensor product of the signatures of the constituent parts. More precisely, if $x : [a, b] \rightarrow V$ and $y : [b, c] \rightarrow V$ are two continuous paths, then their concatenation $x * y$ is the path $[a, c] \rightarrow V$ defined by

$$(x * y)_t := \begin{cases} x_t & \text{if } a \leq t \leq b \\ x_b - y_b + y_t & \text{if } b < t \leq c. \end{cases} \quad (2.30)$$

It then follows from (2.29) that

$$S_{a,c}(x * y) = S_{a,b}(x) \otimes S_{b,c}(y). \quad (2.31)$$

Given any positive integer $n \in \mathbb{Z}_{\geq 1}$ we define *truncated signature* $S_{a,b}^{(n)}(x)$ by

$$S_{a,b}^{(n)}(x) := \Pi_n(S_{a,b}(x)). \quad (2.32)$$

Given a word $\mathbf{K} = (k_1, \dots, k_n) \in \{1, \dots, d\}^n$ we define the *coordinate iterated integral* $S_{a,b}(x)^{\mathbf{K}}$ by

$$S_{a,b}(x)^{\mathbf{K}} := \langle S_{a,b}(x), v_{\mathbf{K}} \rangle_{T(V)} = v_{\mathbf{K}}^*(S_{a,b}^n(x)). \quad (2.33)$$

The name is justified by observing that in the case that the coefficients of the signature are given by (2.28), we see via (2.25) that

$$S_{a,b}(x)^{\mathbf{K}} = \int_{a \leq t_1 \leq \dots \leq t_n \leq b} \dots \int \langle dx_{t_1}, v_{k_1} \rangle_V \dots \langle dx_{t_n}, v_{k_n} \rangle_V. \quad (2.34)$$

The product of two coordinate iterated integrals $S_{a,b}(x)^{\mathbf{K}}$ and $S_{a,b}(x)^{\mathbf{L}}$ yields a quadratic form. But this quadratic form turns out to coincide with a linear functional on elements within the image of the signature in $T(V)$.

Before elaborating, we must introduce the *Shuffle product* \sqcup defined by

$$(i_1, \dots, i_n) \sqcup (j_1, \dots, j_p) := \sum_{\sigma \in \text{Shuff}(n,p)} (a_{\sigma^{-1}(1)}, \dots, a_{\sigma^{-1}(n+p)}) \quad (2.35)$$

where $(a_1, \dots, a_{n+p}) := (i_1, \dots, i_n, j_1, \dots, j_p)$ and

$$\text{Shuff}(n, p) := \{\sigma \in S_{n+p} : \sigma(1) < \dots < \sigma(n) \text{ and } \sigma(n+1) < \dots < \sigma(n+p)\}. \quad (2.36)$$

For $n, p \in \mathbb{N}$, $\text{Shuff}(n, p)$ lists all the ways that two words, of length n and p respectively, can be combined into a single word, of length $n+p$, whilst preserving the order in which the letters of each original word appear.

Returning to the coordinate iterated integrals $S_{a,b}(x)^{\mathbf{K}}$ and $S_{a,b}(x)^{\mathbf{L}}$, on elements in the image of the signature in $T(V)$ their product is given by the coordinate iterated integral associated to the shuffle product of the words \mathbf{K} and \mathbf{L} (cf. Theorem 2.15 in [LCL04]). That is,

$$S_{a,b}(x)^{\mathbf{K}} S_{a,b}(x)^{\mathbf{L}} = S_{a,b}(x)^{\mathbf{K} \sqcup \mathbf{L}}. \quad (2.37)$$

One consequence of (2.37) is that the coordinate iterated integrals span an algebra in $T(V)$. Another consequence is that the signature is a homomorphism of paths with concatenation into the tensor algebra. Reversing a path produces the inverse tensor, and (2.37) establishes that the range is closed under multiplication. Hence the range of the signature forms a group in the tensor algebra.

Informally, it is helpful to think of the range of the signature map as a special curved space in the tensor algebra. As a result there is a lot of valuable structure. A particularly important map is the logarithm; it is one to one on the group and provides a flat parameterisation of the group in terms of elements of the free Lie series (see Section 2 of [LCL04] for details). If $x : [a, b] \rightarrow V$ is a path with signature

$$S_{a,b}(x) = \bigoplus_{n=0}^{\infty} S_{a,b}^n(x) \quad (2.38)$$

then we know that $S_{a,b}^0(x) \equiv 1$. Hence, if we define

$$\hat{S}_{a,b}(x) := \bigoplus_{n=1}^{\infty} S_{a,b}^n(x), \quad (2.39)$$

we formally find via the power series expansion of $t \mapsto \log(1+t)$ that

$$\log S_{a,b}(x) = \bigoplus_{n=1}^{\infty} \frac{(-1)^{n-1}}{n} \left(\hat{S}_{a,b}(x) \right)^{\otimes n} =: \text{LogSig}_{a,b}(x). \quad (2.40)$$

The series defined in (2.40), which is well-defined, is called the *log signature* of x . It can be shown that taking the exponential of the log signature recovers the signature, so no information is lost. The log signature extracts the same information as the signature, but represents it in a more compact way. Analogously to the definition of the truncated signature in (2.32), given any positive integer $n \in \mathbb{Z}_{\geq 1}$ we define the *truncated log signature* (truncated to depth n) by

$$\text{LogSig}_{a,b}^{(n)}(x) := \Pi_n \left(\text{LogSig}_{a,b}(x) \right). \quad (2.41)$$

3. Signature as a Feature Map

Once a model appropriate method of transforming the discrete data points of Ω_V into a continuous path x has been selected, the signature provides an embedding of the dataset Ω_V into a linear space; namely the Hilbert space that is the completion of the tensor algebra $T(V)$. We will frequently refer to the signature of a stream $X \in \mathcal{S}(V)$, with the understanding that we mean the signature of an appropriately chosen continuous path $x : [a, b] \rightarrow V$. The choice of this suitable path x depends on the modelling assumptions.

Since the signature embeds a data stream into a Banach space it is a feature map. We now justify the claims made in Section 2.3 that the signature is an attractive choice of feature map from a machine learning point of view. The first key property is that the infinite graded sequence of statistics provided by the signature identifies the path in an essentially unique manner.

Theorem 3.1 (Uniqueness of signature, [HL10]). *Assume $x : [a, b] \rightarrow V$ is a continuous path of bounded variation. Then the signature $S_{a,b}(x)$ determines x up to tree-like equivalence.*

The full details of *tree-like* equivalence may be found in [HL10]. Heuristically, the signature determines the path up to sections on which the path exactly retraces itself. Thus the signature uniquely determines paths with at least one monotone component, since this prevents the path from retracing itself. As an examples, the path $t \mapsto (t, x_t)$ is always uniquely determined by its signature.

The second key property is that the signature is invariant under reparameterisation.

Theorem 3.2 (Invariance under time reparameterisation). *Assume $x : [a, b] \rightarrow V$ is a continuous path of bounded variation. Let $\psi : [a, b] \rightarrow [a, b]$ be continuously differentiable, increasing and surjective. Then $S_{a,b}(x) \equiv S_{a,b}(x \circ \psi)$.*

The signature essentially quotients out the infinite-dimensional group of time reparameterisations. This allows the signature to record the order of events *without* needing to record precisely when each event occurs, which is an advantage compared to Fourier transform/wavelets techniques.

Consider, for example, monitoring financial markets for signs of insider trading. Three events that might be of interest are a phone call, a trade and price movement. The order of these three events is critical; the order phone call, trade and then price movement could indicate insider trading. To capture the order using Fourier series or wavelets we must sample each of the three channels representing the three events. The sampling frequency will likely need to be very high to resolve each channel accurately enough to capture the order, leading to the consideration of large quantities of uninformative data. But the signature captures this information in the first two levels irrespective of how it is sampled, see example 3.8 in [LLN16].

The third key property is that the terms of the signature decay in size factorially as the order increases. This ensures that the most significant contributions are given by the lower order terms.

Theorem 3.3 (Factorial Decay; Lemma 2.1.1 in [LCL04]). *Assume $x : [a, b] \rightarrow V$ is a continuous path of bounded variation and denote its signature by $S_{a,t}(x) = \bigoplus_{n=0}^{\infty} S_{a,t}^n(x)$. Then for every $n \in \mathbb{N}$ and every $(s, t) \in \Delta_{[a,b]}$ we have*

$$\|S_{s,t}^n(x)\|_{V^{\otimes n}} \leq \frac{\omega_x(s, t)}{n!} \quad (3.1)$$

where $\omega_x : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ is a continuous, sub-additive function, depending on the path x , and vanishing on the diagonal of $\Delta_{[a,b]}$.

Similar decay estimates remain valid for the terms of the signatures of less regular *rough paths*; see [LCL04].

The fourth key property is that continuous functions of paths are approximately linear on signatures. Thus the signature is, in some sense, a *universal nonlinearity* on paths.

Theorem 3.4 (Universal Nonlinearity; Theorem 2.1 in [BDLLS20]). *Let $\mathcal{V}_1([a, b]; V)$ denote the space of continuous paths $[a, b] \rightarrow V$ of bounded variation. Suppose $\mathcal{K} \subset \mathcal{V}_1([a, b]; V)$ is compact and $F : \mathcal{K} \rightarrow \mathbb{R}$ is continuous. Then for any $\varepsilon > 0$ there exists a truncation level $n \in \mathbb{N}$ and, for every $i \in \{0, \dots, n\}$ and any $\mathbf{J} \in \{1, \dots, d\}^i$, coefficients $\alpha_i(\mathbf{J}) \in \mathbb{R}$ such that for every $\theta \in \mathcal{K}$ we have*

$$\left| F(\theta) - \sum_{i=0}^n \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{a,b}(\theta)^{\mathbf{J}} \right| \leq \varepsilon. \quad (3.2)$$

This result can be proven by using the shuffle product relation (2.37) to show that the coordinate iterated integrals span an algebra, before appealing to the *Stone-Weierstrass Theorem*.

Strictly speaking, we must first introduce a topology on $\mathcal{V}_1([a, b]; V)$ before talking about compact subsets. There are, of course, many choices of topology for $\mathcal{V}_1([a, b]; V)$. Our interests lie in the response of various systems to paths, and so it makes sense to use a topology related to this. That is, we would like to induce a topology via a distance metric on $\mathcal{V}_1([a, b]; V)$ satisfying the property that if two paths $\gamma_1, \gamma_2 \in \mathcal{V}_1([a, b]; V)$ are close with respect to this distance, then the response of a system to γ_1 is close to its response to γ_2 .

Defining such a distance metric is a rather challenging task since paths in $\mathcal{V}_1([a, b]; V)$ can have very high oscillation. We make no effort to tackle this rigorously, and instead simply remark that a consequence of *Rough Path Theory* is that such a distance metric exists. The interested reader may find the details in [LCL04], for example. The most important consequence for our purposes is that we can equip $\mathcal{V}_1([a, b]; V)$ with a distance metric, and conclude that if two paths γ_1 and γ_2 are close with respect to this distance, then their signatures $S(\gamma_1)$ and $S(\gamma_2)$ are quantifiably close in $T(V)$. We will always implicitly assume that $\mathcal{V}_1([a, b]; V)$ is equipped with this topology.

We now outline a basic strategy for incorporating the signature as a feature map in a regression task. Recall, we have a Banach space V , a dataset $\Omega_V = \{v_1, \dots, v_N\} \subset V$ for some $N \in \mathbb{N}$, corresponding labels $y_1, \dots, y_N \in \mathbb{R}$, and we want to learn a continuous function $V \rightarrow \mathbb{R}$ using the known responses for the elements in Ω_V .

- Make a model-appropriate choice of embedding of Ω_V in to a path $\gamma : [0, 1] \rightarrow V$ so that $\gamma_{t_k} = v_k$ for all $k \in \{1, \dots, N\}$ for a sequence of times $0 \leq t_1 \leq \dots \leq t_N \leq 1$, and $\gamma \in \mathcal{V}_1([0, 1]; V)$.
- Consider the subset $\mathcal{K}_\gamma \subset \mathcal{V}_1([0, 1]; V)$ defined by

$$\mathcal{K}_\gamma := \{\hat{\gamma}(t_k) : k \in \{1, \dots, N\}\} \quad (3.3)$$

where for each $k \in \{1, \dots, N\}$ the path $\hat{\gamma}(t_k) : [0, 1] \rightarrow V$ is given by $\gamma|_{[0, t_k]}$, i.e. $\hat{\gamma}(t_k)_u := \gamma_{ut_k}$ for $u \in [0, 1]$.

- Choose $K_0 \in \mathbb{N}$ and solve the system of equations

$$\sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{0,1}(\hat{\gamma}(t_k))^{\mathbf{J}} = \sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{0,t_k}(\gamma)^{\mathbf{J}} = y_k \quad (3.4)$$

for each $k \in \{1, \dots, N\}$.

- Take the learnt function to be $F : V \rightarrow \mathbb{R}$ defined by

$$F(v) := \sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{0,1}(\gamma(v))^{\mathbf{J}}, \quad (3.5)$$

where $\gamma(v) : [0, 1] \rightarrow V$ is a model-appropriate alteration of γ to satisfy $\gamma(v)_1 = v$.

The subset $\mathcal{K}_\gamma \subset \mathcal{V}_1([0, 1]; V)$ is compact in the topology induced by the rough path metric we equip $\mathcal{V}_1([0, 1]; V)$ with. Consequently, after composition with the signature, the restriction of our target continuous function to \mathcal{K}_γ is captured by the coordinate iterated integrals (cf. Theorem 3.4); i.e. given any threshold $\varepsilon > 0$, provided a sufficiently large K_0 is chosen, it is possible to solve the system of equations (3.4) up to an error of magnitude at most ε . That is, if we let $\psi : V \rightarrow \mathbb{R}$ denote the unknown continuous function governing the system of interest (i.e. the function we seek to learn), we will be able to find a linear combination of the coordinate iterated integrals of order at most K_0 such that the function $F : V \rightarrow \mathbb{R}$ defined in (3.5) satisfies $|\psi - F| \leq \varepsilon$ throughout Ω_V .

As mentioned previously (cf. Section 2.1), we are implicitly assuming that Ω_V well-approximates V , in some unspecified sense, so that it is reasonable to use the knowledge of ψ on Ω_V to predict its behaviour throughout V . Without making any such assumption, then the performance of a learnt approximation of ψ on new unseen data is referred to as the models *generalisability* (or sometimes its *robustness*). We briefly illustrate particular conditions under which the learnt function F , defined in (3.5), is guaranteed to enjoy good generalisation to certain points *outside* the dataset Ω_V .

A priori there is no guarantee that F gives a good approximation to ψ at *any* point $v \in V \setminus \Omega_V$. Indeed we only know the systems response to the inputs in Ω_V , and without additional assumptions we do not know if this knowledge is sufficient to determine the behaviour of ψ at *any* point outside Ω_V . However, we can use the triangle inequality to make the following observation. Suppose that $v \in V$ and that $\gamma(v) : [0, 1] \rightarrow V$ is a choice of a model-appropriate alteration of γ to satisfy that $\gamma(v)_1 = v$. Then for any $k \in \{1, \dots, N\}$ we have that

$$|\psi(v) - F(v)| \leq |\psi(v) - \psi(v_k)| + |\psi(v_k) - F(v_k)| + |F(v_k) - F(v)| =: \mathbf{A} + \mathbf{B} + \mathbf{C}. \quad (3.6)$$

The magnitude of term \mathbf{A} depends on the regularity of ψ and the distance in V from v to v_k . A consequence of $v_k \in \Omega_V$ is that $\mathbf{B} \leq \varepsilon$. Further, via (3.5) we see that

$$\mathbf{C} \leq \sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} |\alpha_i(\mathbf{J})| \left| S_{0,1}(\gamma(v))^{\mathbf{J}} - S_{0,1}(\hat{\gamma}(t_k))^{\mathbf{J}} \right|. \quad (3.7)$$

Recalling that two paths being close in our choice of *rough path metric* ensures that their signatures are quantifiably close in $T(V)$, we see that (3.7) tells us that the term \mathbf{C} is small provided the paths $\gamma(v)$ and $\hat{\gamma}(t_k)$ are sufficiently close in this *rough path distance*. It therefore follows from (3.6) that we are guaranteed that F and ψ remain close at points $v \in V$ for which the following is true; there exists a $k \in \{1, \dots, N\}$ such that the systems response to v is close to its response to v_k (i.e. that $|\psi(v) - \psi(v_k)|$ is small) and that the paths $\gamma(v), \hat{\gamma}(t_k) \in \mathcal{V}_1([0, 1]; V)$ are sufficiently close in this *rough path distance* sense.

The signature summarises a path in a way that is essentially unique (cf. Theorem 3.1) and that is invariant under reparameterisation of the path (cf. Theorem 3.2). The coefficients of the signature (namely, the coordinate iterated integrals) provide a set of feature functions that have a natural grading (cf. Theorem 3.3), and which are rich enough to capture continuous functions on path space (cf. Theorem 3.4).

The log signature is also an appealing candidate feature map. As previously remarked, the log signature captures the same information as the signature, and represents it in a more compact form. This seems ideal from our machine learning perspective. However, whilst no information about the path is lost, information regarding the nonlinearity is lost.

The coefficients of the signature are rich enough to capture continuous functions on path space via linear combinations (cf. Theorem 3.4). Thus the signature captures the universal nonlinearity of the path, reducing all continuous nonlinearities to linear combinations of its coefficients. This property is not true for the log signature; whilst the coefficients of the log signature can be used to capture continuous functions, it is *not* achieved via linear combinations. Instead, more complicated nonlinear combinations of the log signature coefficients are required. Loosely, the signature captures both the information *and* the nonlinearity of a path, whilst the log signature captures only the information.

4. Controlled Differential Equations and the Log-ODE Method

The *log-ODE* method, developed within the area of *rough path theory* [Lyo98, LCL04, FV10], is a numerical method for solving *Controlled Differential Equations* (CDEs) using the log signature of the control path over short time intervals, rather than relying on pointwise evaluations of the control path. CDEs provide a framework under which one can give a precise meaning for the response of a system to a path according to some specified dynamics.

We first precisely define the notion of a CDE driven by a path. Let $t_0, t_1 \in \mathbb{R}$ with $t_0 < t_1$ and V, W both be a Banach space. Suppose that $X : [t_0, t_1] \rightarrow V$ is a path. Let $\mathbf{L}(V, W)$ denote the set of continuous linear mappings $V \rightarrow W$ and suppose that $f : W \rightarrow \mathbf{L}(V, W)$ is continuous. Then a path $z : [t_0, t_1] \rightarrow W$ solves a CDE controlled, or driven, by X with initial condition $w \in W$ if

$$z_{t_0} = w \quad \text{and} \quad dz_t = f(z_t)dX_t \quad (4.1)$$

for every $t \in (t_0, t_1]$. The CDE framework provides a precise meaning for the response of a system z to the path X according to the dynamics f . In the case that $V := \mathbb{R}$ and $X(t) := t$ for every $t \in [t_0, t_1]$ the CDE (4.1) reduces to an *Ordinary Differential Equation* (ODE). An excellent introduction to CDEs may be found

in [LCL04]. A comprehensive exposition is provided in the textbook [FV10].

The link with signatures is apparent from our earlier observation in Section 2.4 (cf. (2.27)) that the signature of a path X solves the universal non-commuting exponential CDE driven by X . That is, the path $z : [t_0, t_1] \rightarrow T(V)$ defined by $z_t := S_{t_0,t}(X)$ for $t \in [t_0, t_1]$ satisfies that

$$z_{t_0} = \mathbb{1} = 1 \oplus 0 \oplus 0 \oplus \dots \in T(V) \quad \text{and} \quad dz_t = z_t \otimes dX_t \quad (4.2)$$

for every $t \in (0, t_1]$. It turns out that the solution of any linear CDE can be expressed in terms of signature. For example, if the control path $X \in \mathcal{V}_1([t_0, t_1]; V)$ and $f(z_t)dX_t = B(dX_t)z_t$ for a bounded linear map $B : V \rightarrow \mathbf{L}(W, W)$, then the solution to (4.1) is given by

$$z_t = \left(\sum_{n=0}^{\infty} B^{\otimes n} (S_{t_0,t}^n(X)) \right) [w] \quad (4.3)$$

for any $t \in [t_0, t_1]$. This follows from the standard Picard iteration argument coupled with the factorial decay rate for the signature components $S_{t_0,t}^n(X)$ (cf. Theorem 3.3 in this article); see [LCL04, Lyo14], for example. Existence and uniqueness results for the general case using *rough path theory* may be found, for example, in [Lyo98, LCL04, FV10].

The map f in (4.1) can be thought of as a linear map from V to the space of vector fields on W . This object takes an element $v \in V$ and an element $w \in W$ and produces a second element in W representing the infinitesimal change to the state z of the system that will occur if X is changed infinitesimally in the direction v [Lyo14]. The linkage between truncated log signatures and vector fields (see Section 6 in [Lyo14], for example) can be turned into a practical approach to understanding CDEs.

The naive approach of using a Taylor series based method to approximate the solution using truncated log signatures suffer stability issues. The log-ODE method avoids these issues by returning the numerics issues back to state-of-the-art ODE solvers. This is achieved by using the truncated log-signature to determine a path-dependent vector field \tilde{f} such that the solution to the ODE $\frac{dx_t}{dt} = \tilde{f}(x_t)$ provides a good approximation to the solution of the original CDE. Stability issues are now determined by the choice of ODE solver.

The log-ODE method has proven useful in several rapidly developing areas. It is central to the development of *Neural Rough Differential Equations* in [FKLMS21] and in [FLO20] it is illustrated that the log-ODE method may be used for the numerical approximation of *Stochastic Differential Equations* (SDEs). Neural Rough Differential Equations are covered in Section 9 of this article. The log-ODE method proposed in [FLO20] exhibits high order convergence rates that are comparable with all contemporary high order methods for the numerical approximation of SDEs, and can significantly outperform lower order numerical methods for SDEs (such as the Euler-Maruyama and Milstein methods).

We finish this section by following the presentation in Appendix A of [FKLMS21] to illustrate the log-ODE method under the assumption that both V and W are finite dimensional. Treatment of the infinite dimensional case may be found in [BGLY14], for example. We assume $X \in \mathcal{V}_1([0, T]; V)$ has finite length and that $f : W \rightarrow \mathbf{L}(V, W)$ is either linear or bounded with N bounded derivatives for some given $N \in \mathbb{Z}_{\geq 0}$. For $k \in \{0, \dots, N\}$ we define $f^{(k)} : W \rightarrow \mathbf{L}(V^{\otimes k}, W)$ using the derivatives of f as follows (see Definition A.6 in [FKLMS21]). For $k = 0$ we define $f^{(0)}$ by setting $f^{(0)}(y) := y$ for every $y \in W$, and for $k = 1$ we define $f^{(1)}$ by setting $f^{(1)}(y) := f(y)$ for every $y \in W$. For the remaining $k \in \{2, \dots, N\}$ we proceed inductively. Once the function $f^{(k-1)} : W \rightarrow \mathbf{L}(V^{\otimes(k-1)}, W)$ has been defined, we define $f^{(k)} : W \rightarrow \mathbf{L}(V^{\otimes k}, W)$ as follows. For each $y \in W$ we set $f^{(k)}(y)$ to be the unique k -linear map in $\mathbf{L}(V^{\otimes k}, W)$ for which $f^{(k)}(y)[v \otimes v_{k-1}] = Df^{(k-1)}(y)(f(y)[v])[v_{k-1}]$ whenever $v \in V$ and $v_{k-1} \in V^{\otimes(k-1)}$. Here we use the notation that $Df^{(k-1)} : W \rightarrow \mathbf{L}(W, \mathbf{L}(V^{\otimes(k-1)}, W))$ denotes the Fréchet derivative of $f^{(k-1)}$. We additionally use the convention that rounded brackets (\cdot) are used for evaluation at points in W , whilst square brackets $[\cdot]$ have been used for evaluation at points in V and its tensor powers. As remarked in [FKLMS21] the functions $f^{(0)}, \dots, f^{(N)}$ naturally arise in the Taylor expansion associated to the CDE (4.1). A combination of the functions $f^{(0)}, \dots, f^{(N)}$ and the truncated log signature $\text{LogSig}^N(X)$ is used to give an appropriate vector field \tilde{f} .

Suppose that $t_0 \leq s < t \leq t_1$ and that we know the solution at time s , i.e. that we are given z_s . Further suppose that we have compute the truncated log signature $\text{LogSig}_{s,t}^{(N)}(X)$. Then we define a vector

field $\tilde{f} : W \rightarrow \mathbf{L}(V, W)$ by

$$\tilde{f}(u) := \sum_{k=0}^N f^{\circ(k)}(u) \left[\pi_k \left(\text{LogSig}_{s,t}^{(N)}(X) \right) \right] \quad (4.4)$$

for $u \in W$. Here $\pi_k : T^{(N)}(V) \rightarrow V^{\otimes k}$ denotes the projection map. We may then consider the ODE

$$x(0) = z_s \quad \text{and} \quad \frac{dx(r)}{dr} = \tilde{f}(x(r)) \quad (4.5)$$

for all times $r \in (0, 1]$. Then the log-ODE approximation of z_t (given z_s and $\text{LogSig}_{s,t}^{(N)}(X)$) is

$$\text{LogODE}(z_s, f, \text{LogSig}_{s,t}^{(N)}(X)) := x(1). \quad (4.6)$$

In fact the function $\text{LogODE}(z_s, f, \text{LogSig}_{s,t}^{(N)}(X), \cdot) : [s, t] \rightarrow W$ defined by

$$\text{LogODE}(z_s, f, \text{LogSig}_{s,t}^{(N)}(X), \tau) := x\left(\frac{\tau - s}{t - s}\right) \quad (4.7)$$

for $\tau \in [s, t]$ gives an approximation to the solution z_τ for $\tau \in [s, t]$. Only the vector field f and its (iterated) Lie brackets are required to construct the log-ODE vector field \tilde{f} defined in (4.4) (see Remark A.10 in [FKLMS21]). This is a consequence of the fact that the log signature of X lies in a certain free Lie algebra; see Section 2.2.4 of [LCL04] for the precise details.

Error estimates for the log-ODE method may be found in [BGLY14] (including the full general case) where the approximation error of the log-ODE method is quantified in terms of the regularity of the systems vector field f and control path X (also see Appendix B in [FKLMS21] for coverage of the theory in a simplified setting). These results make use of numerous technical details from *rough path theory*, which one may find in either of [LCL04, FV10], for example. The simplified conclusion of the sophisticated theory is that log-ODEs can approximate CDEs. Moreover, the convergence rate is controlled by the depth of log signature used and the step size considered (i.e. the size of $t - s$ in our presentation above). Performance of the log-ODE method can be improved by appropriately selecting both the step size and the truncation depth [BGLY14].

5. Computing Signatures

Making use of signature based techniques requires being able to explicitly compute signatures of paths, or rather being able to write code to enable a computer to carry out the computations. Whilst there are now several packages for computing path signatures (the *iisignature* package [GB18] and the *Signatory* package [KL20] for example), we will focus on the first dedicated package created for this purpose, namely the *esig* package [Lyo02]. The *esig* package has been developed by Terry Lyons et al. since 2002 with numerous refinements and updates being made, including extensions implemented by three summer interns at the Alan Turing institute in 2017.

The *esig* package provides implementations of a series of mathematical tools for handling the signature of a path. It is a python package powered by the underlying *libalgebra* C++ package. The documentation for *esig* may be found at <https://esig.readthedocs.io/en/latest/index.html>, and may be installed directly via **pip** (a standard package-management system used to install and manage software packages written in Python). It speedily computes truncated signatures and log signatures of paths that are well-approximated by piecewise linear functions.

Observe that if $\gamma : [a, b] \rightarrow V$ is linear, so that $\gamma_t = \gamma_a + \frac{t-a}{b-a}(\gamma_b - \gamma_a)$ for $t \in [a, b]$, then

$$S_{a,b}(\gamma) = \bigoplus_{n=0}^{\infty} \frac{1}{n!} (\gamma_b - \gamma_a)^{\otimes n} =: \exp(\gamma_b - \gamma_a). \quad (5.1)$$

Therefore, if we have a piecewise linear path $\gamma : [0, T] \rightarrow V$ connecting the points $x_1, \dots, x_k \in V$, we can repeatedly use Chen's identity (specifically (2.31)) and (5.1) to deduce that

$$S(x_1, \dots, x_k) := S_{0,T}(\gamma) = \exp(x_2 - x_1) \otimes \exp(x_3 - x_2) \otimes \dots \otimes \exp(x_k - x_{k-1}). \quad (5.2)$$

Hence computing the truncated signature $S_{0,T}^{(N)}(\gamma) := \Pi_N(S_{0,T}(\gamma))$ only requires evaluating the terms up to order N on the RHS of (5.2).

In the example above, the array of points x_1, \dots, x_k contains a full description of the path γ . The `esig` package takes such an array as argument, and computes the signature (or logsignature) of the associated path truncated at a chosen degree. Computational speed is achieved via efficient algorithms making use of the decomposition of the signature illustrated in (5.2).

The `esig` package is a python package, and the function `stream2sig` requires an input array that is a `numpy.array((l, d), dtype=float)` of shape (l, d) where l is the number of sample points in the path (including both the start and end points) and d is the dimension of the target space in which the points live. As an example, suppose we wanted to compute the signature of an affine path $\gamma : [a, b] \rightarrow \mathbb{R}^d$. Then the following python code computes the signature of γ truncated to depth N .

```
>>> import numpy as np
>>> import esig
>>> x = [float( $\gamma_a$ ), float( $\gamma_b$ )]
>>> path = np.array(x)
>>> sig = esig.stream2sig(path, N)
```

Table 1: Python code to compute $S_{a,b}^{(N)}(\gamma)$ using `esig`

The input stream needs to be a numpy array of floats; input numpy arrays of other dtype will not work. The output `sig` is a one-dimensional array of numbers that is interpreted as the coefficients of the basis elements produced by `sigkeys` delivered in the same order. The function `sigkeys` takes two integers (m, k) as input, and returns a basis of the tensor algebra on an alphabet of size m up to depth k as a text string.

```
>>> esig.sigkeys(2,3)
'(1) (2) (1,1) (1,2) (2,1) (2,2) (1,1,1) (1,1,2) (1,2,1) (1,2,2) (2,1,1) (2,1,2) (2,2,1) (2,2,2)'
```

Table 2: Basis of tensor algebra on an alphabet of size 2 up to depth 3 enumerated by `sigkeys`

Hence, if $d = 2$ and $N = 3$ in the code in Table 1, the entry `sig[5]` is the value of the coordinated iterated integral of the path corresponding to the word $(2, 1)$. As expected, the string returned in Table 2 consists of $\sum_{i=0}^3 2^i = 15$ terms.

For a concrete example, suppose $\gamma : [0, 1] \rightarrow \mathbb{R}^3$ is given by $\gamma(t) := (3t + 2, t - 2)$ and we want to compute $S_{0,1}(\gamma)$ truncated to order 3. This can be accomplished via the following python code.

```
>>> import numpy as np
>>> import esig
>>> x = [[2.0, -2.0], [5.0, -1.0, ]]
>>> path = np.array(x)
>>> sig = esig.stream2sig(path, 3)
```

Table 3: Code to compute $S_{0,1}^{(3)}(\gamma)$

The output of the code in Table 3 is

```
array([ 1.    , 3.    , 1.    , 4.    , 1.5 ,
        1.5   , 0.5   , 4.5   , 1.5   , 1.5
        0.5   , 1.5   , 0.5   , 0.5   , 0.16666667 ])
```

Table 4: Output of python code in Table 3

Following the enumeration of the basis illustrated in Table 2, this lists the values of the various coordinated iterated integrals for all the words in two letters up to a maximum length of three. In particular, the coordinate iterated integral of γ corresponding to the word $(2, 1)$ is 1.5.

We can now give a toy problem illustrating the signatures ability to determine the order of events. For this purpose, we consider a path γ in \mathbb{R}^2 that starts from the origin $(0, 0)$. We consider a total of 30000 samples from this path. For a random integer m between 1 and 29998, the path will remain at $(0, 0)$ for the first m samples. Over the next two samples, it moves to $(1, 1)$, with an equal chance of following the route $(0, 0) \rightarrow (1, 0) \rightarrow (1, 1)$ or the route $(0, 0) \rightarrow (0, 1) \rightarrow (1, 1)$. After reaching the point $(1, 1)$, the path remains here for the remainder of the 30000 samples. Our task is to determine which route the path followed; namely, did it first move away from 0 in the x coordinate or the y coordinate?

This is easily determined by computing the signature truncated to depth 3, as illustrated in Figure 8.

```
array([ 1.00000000e+00, 1.00000000e+00, 1.00000000e+00, 5.00000000e-01,
        1.00000000e+00, 0.00000000e+00, 5.00000000e-01, 1.66666667e-01,
        5.00000000e-01, -2.77555756e-17, 5.00000000e-01, -1.38777878e-17,
        0.00000000e+00, -1.38777878e-17, 1.66666667e-01 ])
```

(a) Esig output for signature truncated to depth 3 when γ follows the route $(0, 0) \rightarrow (1, 0) \rightarrow (1, 1)$

```
array([ 1.00000000e+00, 1.00000000e+00, 1.00000000e+00, 5.00000000e-01,
        1.00000000e+00, 0.00000000e+00, 5.00000000e-01, 1.66666667e-01,
        -1.38777878e-17, 0.00000000e+00, -1.38777878e-17, 5.00000000e-01,
        -2.77555756e-17, 5.00000000e-01, 1.66666667e-01 ])
```

(b) Esig output for signature truncated to depth 3 when γ follows the route $(0, 0) \rightarrow (0, 1) \rightarrow (1, 1)$ Figure 8: Signatures of path γ truncated to depth 3 for the two possible routes.

The returned numbers are the coefficients of the coordinate iterated integrals corresponding to the words in the order dictated by `sigkeys(2,3)` (illustrated in Table 2). Using the basis expansion given by `sigkeys(2,3)`, we see that differences appear in the depth 3 terms. These differences capture the order in which the coordinates change from 0 to 1. For example, in both Figure 8 (a) and (b) the 9th entry in each array corresponds to the coordinate iterated integral $(1, 1, 2)$. The positive sign of this entry in (a) corresponded with the change in the dimension corresponding to the word '1' (i.e. the x coordinate) happens before the change in the dimension corresponding to the word '2' (i.e. the y coordinate). The negative sign of this entry in (b) corresponds with the order of the changes being reversed.

The following python code computes the signature truncated to depth N of the piecewise linear path visiting the points $x_1, \dots, x_m \in \mathbb{R}^d$ in this order.

```
>>> import numpy as np
>>> import esig
>>> x = [x1, ..., xm]
>>> path = np.array(x)
>>> sig = esig.stream2sig(path, N)
```

Table 5: Code to compute truncated signature of order N of the piecewise linear path visiting the points $x_1, \dots, x_m \in \mathbb{R}^d$ in this order

In an analogous manner, the function `esig.stream2logsig` can compute the log signature of a piecewise linear path visiting a given number of points. It again returns a string of coefficients, though the basis with which they

correspond is more complicated. More detailed explanations and examples can be found in the documentation for `esig`.

It is both informative and instructive to work through some examples of using signatures within a relatively simple machine learning context. There are a number of demonstration notebooks available via the DataSig website. The *Introduction to Path Signatures* python notebook is designed to familiarise the user with the key concepts. Primarily aimed at data scientists and machine learning practitioners, it uses simple illustrative examples to explore the behaviour of path signatures and of the affect that various stream transformations (including those covered in Section 2.2) have on the signature of a stream.

The *Handwritten Digit Classification* python notebook demonstrates the use of path signatures (via `esig`) for handwritten digit classification. Given sequences of pen strokes contained in the MNIST sequence dataset, corresponding to the MNIST handwritten digit data set by Edwin D. De Jong, path signatures of transformations of these streams are computed and incorporated as features into a linear classifier. The notebook additionally explores combining the features with unsupervised classification techniques.

6. Expected Signature

The signature provides a summary of a path that is adequate for predicting its effect on a broad range of different systems whilst managing to discard irrelevant information. Frequently we are interested in understanding more than a single path at a time; that is, we want to adequately summarise a collection of paths in order to predict their collective effect on systems. The paths traced out by multiple points on an object will give more information about the objects trajectory and motion than a single point. The effect of Covid-19 restrictions on the reproduction (R) number of the virus are more accurately predicted knowing the levels of compliance of numerous people across several locations, as opposed to relying on the level of compliance of a single individual.

A distribution can summarise a collection of paths; for example, a finite collection $\mathcal{N} \subset \mathcal{V}_1([a, b]; V)$ can be summarised by the empirical measure

$$\delta_{\mathcal{N}} := \frac{1}{|\mathcal{N}|} \sum_{\gamma \in \mathcal{N}} \delta_{\gamma}. \quad (6.1)$$

We would like to summarise distributions in an analogous manner to how the signature summarises paths themselves. That is, discarding irrelevant information whilst retaining enough detail to distinguish between different distributions. Usually this requires restricting the class of distribution considered by imposing additional constraints (cf. Brownian Motion, Markov processes etc). It turns out that the signature allows us to summarise *all* distributions without restriction.

If we let $\mathcal{PV}_1([a, b]; V)$ denote the space of (Borel) probability measures on $\mathcal{V}_1([a, b]; V)$ then the *Expected Signature* is the map $\mathbb{S} : \mathcal{PV}_1([a, b]; V) \rightarrow T(V)$ defined by

$$\mathbb{S}(\mu) := \mathbb{E}_{\mu} [S_{a,b}(X)] = \bigoplus_{n=0}^{\infty} \mathbb{E}_{\mu} [S_{a,b}^n(X)]. \quad (6.2)$$

Given a stochastic process X_t on V for $t \in [a, b]$ under a probability space $(\mathcal{N}, \mathbb{P})$, we can push-forward the measure \mathbb{P} to $\mathcal{V}_1([a, b]; V)$ via the map $Q : \mathcal{N} \rightarrow \mathcal{V}_1([a, b]; V)$ defined by $Q(\omega) := (s \mapsto X_s(\omega))_{s \in [a, b]}$. The resulting measure $Q_*\mathbb{P}$ is an element of $\mathcal{PV}_1([a, b]; V)$, thus we can make sense of the *expected signature* of X_t as $\mathbb{S}(Q_*\mathbb{P})$, with it being common to abuse notation and write $\mathbb{S}(X)$. This is the natural generalisation of the moments of the process X_t , and moreover, in [BO19], it is shown that the *Log Expected Signature* $\log \mathbb{S}(X)$ provides the natural generalisation of the cumulants of X_t .

Beginning with the work of Fawcett in his thesis [Faw03], there has been a large number of works investigating the relationship between the expected signature $\mathbb{S}(X)$ and the law of $S(X)$. Fawcett originally proved that if the measure $Q_*\mathbb{P}$ is compactly supported, then the law of $S(X)$ is uniquely determined by $\mathbb{S}(X)$. Subsequently, it has been established that if the radius of convergence of the power series $\sum_{n=0}^{\infty} z^n \mathbb{E} \|S^n(X)\|$ is infinite, then the expected signature uniquely determines the law of $S(X)$ [Che13]. Further refinements to

the extent to which the expected signature of a stochastic process X_t determines the law of X_t may be found in [CL13, CO18].

For Brownian motion with Lévy area B_t on a bounded C^1 domain $\Omega \subset \mathbb{R}^d$, the investigation of the finiteness of the radius of convergence for the power series $\sum_{n=0}^{\infty} z^n \mathbb{E} \|S^n(B_t)\|$ was initiated in [LN11], where the authors established that the radius of convergence was strictly positive. In [CL13] it is proven that this radius of convergence is infinite, provided B_t is considered up to a fixed finite time $T > 0$. Consequently, for such Brownian motions, the expected signature determines the law. But recently it has been established that this is *not* true for the exit time of Brownian motion in two-dimensions. More concretely, if B_t^z is the standard Brownian motion in \mathbb{R}^2 started at $z \in \mathbb{B}^2(0, 1)$ and stopped at the first time it hits the boundary $\partial \mathbb{B}^2(0, 1)$, then radius of convergence of the corresponding power series is finite; see [BDMN19] for full details.

Of greater interest from the machine learning perspective is that the expected signature provides a systematic way to describe probability measures on paths in terms of their effects. An early breakthrough was the realisation that complex path measures (such as the Wiener measure) can be effectively approximated by a measure supported on finitely many paths having the same expected signature on a truncated tensor algebra $T^{(n)}(V)$ [LV04, LL11, LY13]. Unlike the *Martingale Problem* approach originating in [SV79], there is no restriction on the probability measures considered.

A result from [CO18] establishes that the expected signature essentially uniquely determines distributions on compact subsets of $\mathcal{PV}_1([a, b]; V)$. It is proven in Theorem 5.6 of [CO18] that when restricted to compact subsets, the expected signature is injective up-to tree-like equivalence (see Appendix B of [CO18] for full details). The important point from our perspective is that if the set of tree-like paths is removed from $\mathcal{V}_1([a, b]; V)$, then the expected signature is injective on compact subsets of $\mathcal{PV}_1([a, b]; V)$.

The work [BDLLS20] establishes a universality property for the expected signature analogous to the universality property of the signature itself (cf. Theorem 3.4). For the signature, this universality is that, on compact subsets of path-space, the components of the signature (i.e. the coordinate iterated integrals) span the set of continuous functions. An almost identical result turns out to be true for the space $C^0(\mathcal{K})$ with $\mathcal{K} \subset \mathcal{V}_1([a, b]; V)$ compact, provided we first lift the expected signature map so that it maps a probability measure μ to a path $[a, b] \rightarrow T(V)$, rather than a single element of $T(V)$.

This lift, called the *Pathwise Expected Signature*, is the map $\Phi_{\text{Path}} : \mathcal{PV}_1([a, b]; V) \rightarrow \mathcal{V}_1([a, b]; T(V))$ defined by

$$\Phi_{\text{Path}}(\mu)_t := \mathbb{E}_{\mu} [S_{a,t}(X)] \quad \text{so that} \quad \Phi_{\text{Path}}(\mu)_b = \mathbb{S}(\mu). \quad (6.3)$$

The coordinate iterated integrals forming the components of $S_{a,b}(\Phi_{\text{Path}}(\mu))$ then span $C^0(\mathcal{K})$ for compact subsets $\mathcal{K} \subset \mathcal{V}_1([a, b]; V)$.

Theorem 6.1 (Expected Signature Universality; Theorem 3.2 in [BDLLS20]). *Let $\mathcal{C}([a, b]; V) \subset \mathcal{V}_1([a, b]; V)$ denote the subset resulting from removing all tree-like paths. Suppose $\mathcal{K} \subset \mathcal{C}([a, b]; V)$ is compact and consider a continuous function $F : \mathcal{K} \rightarrow \mathbb{R}$. Then for any $\varepsilon > 0$ there exists a truncation level $m \in \mathbb{N}$ and, for every $k \in \{0, \dots, m\}$ and any $\mathbf{J} \in \{1, \dots, d\}^k$, coefficients $\alpha_k(\mathbf{J}) \in \mathbb{R}$ such that for any $\mu \in \mathcal{P}(\mathcal{K})$ we have*

$$\left| F(\mu) - \sum_{k=0}^m \sum_{\mathbf{J} \in \{1, \dots, d\}^k} \alpha_k(\mathbf{J}) S_{a,b}(\Phi_{\text{Path}}(\mu))^{\mathbf{J}} \right| < \varepsilon. \quad (6.4)$$

The components $S_{a,b}^n(\Phi_{\text{Path}}(\mu))$ have the factorial decay associated with all signature components (cf. Theorem 3.3). Consequently, via the same reasoning used for the signature of a path in Section 3, the pathwise expected signature gives a universal feature map providing a graded set of feature functions (the coordinate iterated integrals of $S(\Phi_{\text{Path}}(\mu))$) that are detailed enough to essentially uniquely determine the distribution μ . And once again the decay of the coordinate iterated integrals tells us that the more easily computed lower order terms will typically be more informative than the higher order ones. The pathwise expected signature provides a route to tackling distribution regression. This is elaborated on in Section 13 of this article.

The *Drone Identification* python notebook is an instructive example of using expected path signatures in a classification task that may be worked through by the reader. In the notebook a classifier for identifying drones is constructed based on the following assumption. When a radio pulse is reflected off a drone, the reflected signal received back by the observer is a combination of the reflection caused by the drone's body and

the reflection caused by the drone's propeller. The expected path signature is used to characterise the random behaviour in reflected signals. Estimates of expected path signatures are used as feature vectors for the task of distinguishing between drone and non-drone objects. The task of predicting the number of rotations per minute (rpm) of a drone's propeller is additionally considered.

7. Truncation Order Selection

The signature of a stream lives in the infinite-dimensional tensor algebra $T(V)$. Consequently, applications seemingly require the selection of some finite sub-collection of terms. The factorial decay of the components of the signature (cf. Theorem 3.3) means it is common to choose the first N terms since these will typically be the largest. Recall that the first N terms of a signature $S(x)$ are called the *truncated signature of depth N* and denoted by $S^{(N)}(x) := \Pi_N(S(x))$. Whilst such a truncation captures the largest components, it nevertheless loses information captured by the higher order terms. By using this naive truncation we will be unable to accurately learn *any* functions with non-trivial dependence on these higher order terms.

Introducing a suitable augmentation of the dataset *before* taking the signature offers a solution to this problem. Recall (cf. Section 2.2) that by augmenting the dataset Ω_V , we mean considering a map $\Theta : V \rightarrow W$ for some Banach space W , and working with the augmented dataset $\Theta(\Omega_V)$. If we take $W := V \times U$ and $\Theta(v) := (v, \varphi(v))$ for a suitable $\varphi : V \rightarrow U$, the information from the higher order terms of the signature of the stream Ω_V is contained within the truncated signature of the stream $\Theta(\Omega_V)$. In practice, V is often given by \mathbb{R}^d for some $d \in \mathbb{N}$, and we take $U := \mathbb{R}^e$ for some $e \in \mathbb{N}$ so that $W = \mathbb{R}^{d+e}$.

We have already seen several explicit examples of possible augmentations in Section 2.2, and each augmentation offers differing particular benefits. A Lead-Lag augmentation aides the study of a streams variance via its signature (see [CK16]), the invisibility-reset augmentation (cf. (2.15)) ensures that the signature of the transformed stream will contain information on the initial position which is otherwise lost due to the translation-invariance of the signature, and the time augmentation (cf. (2.11)) introduces a monotonic coordinate that can ensure the resulting path is uniquely determined by its signature (cf. Theorem 3.1).

Many works choose the augmentation map Θ via experimentation with several explicit options, see [CO18, KO19] for example. Recently, it has been proposed in [ABKLS19] that the augmentation map should be data dependent and be learnt, i.e to consider $\Theta = \Theta^\theta$ for a trainable parameter θ . This removes any limitations on the form of Θ , allowing complicated nonlinear influences from the higher order signature terms to be accurately captured. In particular, the augmentation may form part of a *Neural Network*.

Initially inspired by the biological workings of a human brain, a neural network is a collection of functions (called nodes) that are arranged in layers. Each layer connects to its neighbours, and the output of a layer is used as the input data for the next layer. Typically, each node has an associated *activation function* controlling whether or not the nodes function is run depending on the received input. Empirically, simple initial inputs and a large number of layers works better than more complex initial specifications teamed with fewer layers. Mathematically, neural networks may be thought of as linear functions composed with nonlinearity. Whilst complicated nonlinear transformations are allowed between layers, all optimisation happens at the layers themselves, where we are always seeking a linear combination of a specified collection of functions. Detailed and thorough introductions to neural networks may be found in [Gér19], [Hea12], [Gur97], for example.

Recall that, on compact subsets, continuous functions of paths can be arbitrarily well-approximated by linear combinations of coordinate iterated integrals (cf. Theorem 3.4). In this sense, the signature is a universal nonlinearity of streams, which makes it a candidate for the nonlinear transformation between layers within a neural network. Using the signature in this manner requires adapting the signature to map a stream to a stream, rather than mapping a stream to statistics with no obvious stream-like properties. In [ABKLS19] this is achieved by lifting the input stream Ω_V to a stream of streams, making use of the observation that a path $X : [a, b] \rightarrow V$ induces a path-of-paths in path space by mapping $t \mapsto X|_{[a, t]}$ for $t \in [a, b]$.

This idea makes it easy to see that a stream (x_1, \dots, x_k) can induce the stream-of-streams given by $((x_1, x_2), (x_1, x_2, x_3), \dots, (x_1, \dots, x_k))$, for example. But there is no unique way to induce a stream-of-streams from a stream. A generic *lift* $l : \mathcal{S}(V) \rightarrow \mathcal{S}(\mathcal{S}(V))$ can be considered, with the signature of $l(X)$ for $X \in \mathcal{S}(V)$ to be the stream of signatures of the components. More precisely, if $X \in \mathcal{S}(V)$ then $l(X) =$

$(l(X)_1, \dots, l(X)_m)$ for some $m \in \mathbb{N}$, where $l(X)_j \in \mathcal{S}(V)$ for each $j \in \{1, \dots, m\}$. The signature of $l(X)$ is then taken to be the stream $(S(l(X)_1), \dots, S(l(X)_m)) \in \mathcal{S}(T(V))$. The particular choice of lift will be determined by the modelling assumptions. Several example lifts l may be found in [ABKLS19].

Having modified the signature to transform a stream to a stream, the authors define a *deep signature model* in which several learnt augmentation steps are implemented. The output of each layer is a linear function of the signature, and the signature of this forms the input functions for the next layer. The shuffle product plays a vital role in understanding how to compute $S(L(S(X)))$ for a linear function L [LCL04]. Two simple architectures involving signature layers are shown in Figure 9, and full details of the proposed models can be found in [ABKLS19].




Figure 9: Two simple architectures illustrating the possible inclusion of a signature layer - Figure 1 in [ABKLS19]. Reproduced with permission. Copyright held by authors of [ABKLS19].

The required computation and back-propagation of the signature transform is efficiently managed by the *Signatory* package [KL20] (cf. Section 5). The resulting model obtained in [ABKLS19] is compared against several other techniques via the task of learning the *Hurst parameter* H of a fractional Brownian motion B^H . That is, to learn the map $\mathbf{x}^H \mapsto H$ for an implementation

$$\mathbf{x}^H = ((t_0, B_{t_0}^H), \dots, (t_n, B_{t_n}^H)) \in \mathcal{S}(\mathbb{R}^2) \quad (7.1)$$

for some realisation of B^H . Empirical evidence suggests fractional Brownian motions are closely linked with financial market data [GJR18], and estimating the Hurst parameter is a non-trivial task [LLN09].

A total of 700 samples was considered, with 600 being used as the training set and the remaining 100 forming the test set. Each individual sample was an instance of fractional Brownian motion, sampled at 300 time steps of the interval $[0, 1]$, with Hurst parameters in the range $[\frac{1}{5}, \frac{4}{5}]$. Every model was trained for 100 epochs with the loss function take to be the mean squared error (MSE). The results are summarised in the following table.

Model	Mean	Variance
Rescaled Range	7.2×10^{-2}	3.7×10^{-3}
LSTM	4.3×10^{-2}	8.0×10^{-3}
Feedforward	2.8×10^{-2}	3.0×10^{-3}
Neural-Sig	1.1×10^{-2}	8.2×10^{-4}
GRU	3.3×10^{-3}	1.3×10^{-3}
RNN	1.7×10^{-3}	4.9×10^{-4}
DeepSigNet	2.1×10^{-4}	8.7×10^{-5}
DeeperSigNet	1.6×10^{-4}	2.1×10^{-5}

Table 6: Final test MSE for the different models - Variant of Table 1 in [ABKLS19]

The *Rescaled Range* method is a mathematically derived method involving no learning, see [Hur51]. The *Long Short-Term Memory* (LSTM), GRU and *Recurrent Neural Network* (RNN) models provide baselines in the context of recurrent neural networks. The *Neural-Sig* model provides a baseline in the context of signatures, with only a single layer of augmentation being used. The *DeepSigNet* and *DeeperSigNet* are both deep signature models featuring multiple layers of augmentation. Whilst the traditional signature based methods perform

worse than the recurrent models, the deep signature models outperform all other models by at least an order of magnitude. Full details of this comparison and the baseline models may be found in section 4.2 and appendix B.2 of [ABKLS19].

8. Signature Kernel

The signature provides a universal feature map embedding the dataset Ω_V into the tensor algebra $T(V)$. We can naturally equip $T(V)$ with an inner product $\langle \cdot, \cdot \rangle$ (cf. Section 2.4). This allows us to define the *Signature Kernel*

$$K_{x,y}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{T(V)} \quad (8.1)$$

for $x \in \mathcal{V}_1([a, b]; V)$, $y \in \mathcal{V}_1([c, d]; V)$, $s \in [a, b]$ and $t \in [c, d]$.

The completion $\overline{T(V)}$ is a Hilbert space, meaning that the signature provides us with a kernel. If we let $\gamma : [0, 1] \rightarrow V$ denote the continuous path that the dataset Ω_V is transformed to, with $\gamma_{t_k} = x_k$ for each $x_k \in \Omega_V$, then we may consider the associated function $K_{\Omega_V} : \Omega_V \times \Omega_V \rightarrow \mathbb{R}_{\geq 0}$ given by

$$K_{\Omega_V}(x_i, x_j) := K_{\gamma, \gamma}(t_i, t_j) = \langle S_{0,t_i}(\gamma), S_{0,t_j}(\gamma) \rangle_{T(V)} \quad (8.2)$$

for $x_i, x_j \in \Omega_V$. In [CFLSY20] the authors consider the use of the kernel function K_{Ω_V} , defined in (8.2), for regression tasks. Via their observation that the function $(s, t) \mapsto K_{x,y}(s, t)$ defined in (8.1) satisfies a linear second-order hyperbolic PDE (a *Goursat problem* [Lee60]), the authors introduce a method to use the full signature transform for regression tasks *without* any truncation via use of the kernel function K_{Ω_V} . The primary aim for the remainder of this section is to summarise this approach developed in [CFLSY20].

In [KO19], it is shown that a kernel on the space carrying the data (V in our setting) implies a kernel on truncated signatures. This follows from the observation that if we have embeddings $\Psi : V \rightarrow W$ and $\Phi : V \rightarrow W^*$ (i.e. a kernel on V), then one can define a kernel via

$$\Psi_{\text{SIG}}^N : \mathcal{V}_1([0, 1]; V) \rightarrow T^{(N)}(\mathcal{V}_1([0, 1]; W)) \quad \text{defined by} \quad \Psi_{\text{SIG}}^N(\gamma) := S_{0,1}^{(N)}(\Psi \circ \gamma), \quad (8.3)$$

and

$$\Phi_{\text{SIG}}^N : \mathcal{V}_1([0, 1]; V) \rightarrow T^{(N)}(\mathcal{V}_1([0, 1]; W))^* \quad \text{defined by} \quad \Phi_{\text{SIG}}^N(\gamma) := S_{0,1}^{(N)}(\Phi \circ \gamma). \quad (8.4)$$

That Φ_{SIG}^N defined in (8.4) maps into $T^{(N)}(\mathcal{V}_1([0, 1]; W))^*$ as claimed relies on both the observations that $\mathcal{V}_1([0, 1]; W^*) \subset \mathcal{V}_1([0, 1]; W)^*$ and $T^{(N)}(\mathcal{V}_1([0, 1]; W)^*) \subset T^{(N)}(\mathcal{V}_1([0, 1]; W))^*$. In [KO19] it is explained how to use truncated kernels to approximate the full untruncated signature kernel. For paths of bounded variation, there are efficient algorithms to compute this truncated kernel, with the practical consequences explored in [OT20].

In [CFLSY20] the full untruncated signature kernel is studied. It is established that for two continuously differentiable paths x and y , the full untruncated signature kernel $K_{x,y}(\cdot, \cdot)$ solves the following PDE.

Theorem 8.1 (Signature Kernel Solves PDE; Theorem 2.5 in [CFLSY20]). *Let $I = [a, b]$ and $J = [c, d]$ be compact intervals and suppose $x \in C^1(I; V)$ and $y \in C^1(J; V)$ for some Banach space V . Consider the bilinear form $K : I \times J \rightarrow \mathbb{R}$ defined by*

$$K_{x,y}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{T(V)}. \quad (8.5)$$

Then $K_{x,y}(\cdot, \cdot)$ solves the following linear second-order hyperbolic PDE

$$\frac{\partial^2 K_{x,y}}{\partial s \partial t} = \langle \dot{x}_s, \dot{y}_t \rangle K_{x,y} \quad (8.6)$$

with initial conditions $K_{x,y}(a, \cdot) = K_{x,y}(\cdot, c) = 1$ and where $\dot{x}_s := \frac{dx_p}{dp} \Big|_{p=s}$ and $\dot{y}_t := \frac{dy_q}{dq} \Big|_{q=t}$.

The PDE in (8.6) is an example of a *Goursat problem*, for which existence and uniqueness results originate in [Lee60]. Consequently, for sufficiently regular paths, the existence and uniqueness of K solving (8.5) is guaranteed (cf. Theorem 3.1 in [CFLSY20]). The content of Theorem 8.1 makes sense for rougher paths. That is, the PDE (8.6) can be made sense of for paths x and y of lower regularity by considering its integral form

$$K_{x,y}(s, t) = 1 + \int_a^s \int_c^t \langle S_{a,p}(x), S_{c,q}(y) \rangle_{T(V)} \langle dx_p, dy_q \rangle_V \quad (8.7)$$

and making an appropriate definition for the quantity $\langle dx_p, dy_q \rangle_V$. This is done in [CFLSY20] for the broader class of *geometric rough paths*; see section 4 in [CFLSY20] for full details.

Numerical schemes to approximate $K_{x,y}$ are developed in [CFLSY20], which we now illustrate in the case that $V = \mathbb{R}^d$ for some $d \in \mathbb{N}$. Assume that $I, J \subset \mathbb{R}$ are compact intervals and that $x : I \rightarrow \mathbb{R}^d$ and $y : J \rightarrow \mathbb{R}^d$ are both piecewise linear. Then the PDE (8.6) becomes

$$\frac{\partial^2 K_{x,y}}{\partial s \partial t} = C K_{x,y} \quad (8.8)$$

on each domain $\mathcal{D}_{ij} := [u_i, u_{i+1}] \times [v_j, v_{j+1}]$ that $C := \langle \dot{x}_s, \dot{y}_t \rangle$ is constant on. In integral form, (8.8) reads

$$K_{x,y}(s, t) = K_{x,y}(s, v) + K_{x,y}(u, t) - K_{x,y}(u, v) + C \int_u^s \int_v^t K_{x,y}(r, w) dr dw \quad (8.9)$$

for $(s, t), (u, v) \in \mathcal{D}_{ij}$ with $u \leq s$ and $v \leq t$. Approximations of the double-integral in (8.9) can be used to develop finite difference schemes to approximate $K_{x,y}$. The relatively simple approximation

$$\int_u^s \int_v^t K_{x,y}(r, w) dr dw \approx \frac{1}{2} (K_{x,y}(s, v) + K_{x,y}(u, t)) (s - u)(t - v), \quad (8.10)$$

is used in [CFLSY20] to develop such a scheme.

In order to more precisely discuss this finite difference scheme, let I, J be compact intervals and $\mathcal{D}_1 := \{u_0 < u_1 < \dots < u_{m-1} < u_m\}$ and $\mathcal{D}_2 := \{v_0 < v_1 < \dots < v_{n-1} < v_n\}$ be partitions of I and J respectively for which x is piecewise linear with respect to \mathcal{D}_1 and y is piecewise linear with respect to \mathcal{D}_2 . Then (8.10) can be used to define numerical schemes for (8.6) on the grid $P_0 := \mathcal{D}_1 \times \mathcal{D}_2$, and its dyadic refinements. For $\lambda \in \mathbb{N}_0$ we define P_λ as the dyadic refinement of P_0 such that

$$P_\lambda \cap ([u_i, u_{i+1}] \times [v_j, v_{j+1}]) = \bigcup_{k,l=0}^{2^\lambda} \left\{ \left(u_i + \frac{k}{2^\lambda} (u_{i+1} - u_i), v_j + \frac{l}{2^\lambda} (v_{j+1} - v_j) \right) \right\}. \quad (8.11)$$

In [CFLSY20] the authors define the finite difference scheme

$$\begin{aligned} \hat{K}(s_{i+1}, t_{j+1}) &= \hat{K}(s_{i+1}, t_j) + \hat{K}(s_i, t_{j+1}) - \hat{K}(s_i, t_j) + \\ &\quad \frac{1}{2} \langle x_{s_{i+1}} - x_{s_i}, y_{t_{j+1}} - y_{t_j} \rangle \left(\hat{K}(s_{i+1}, t_j) + \hat{K}(s_i, t_{j+1}) \right), \end{aligned} \quad (8.12)$$

with $\hat{K}(s_0, \cdot) = \hat{K}(\cdot, t_0) = 1$, on the grid

$$P_\lambda = \bigcup_{i=0}^{m-1} \bigcup_{j=0}^{n-1} \bigcup_{k,l=0}^{2^\lambda-1} \left\{ \left(u_i + \frac{k}{2^\lambda} (u_{i+1} - u_i), v_j + \frac{l}{2^\lambda} (v_{j+1} - v_j) \right) \right\} =: \bigcup_{i=0}^{2^\lambda n} \bigcup_{j=0}^{2^\lambda m} \{(s_i, t_j)\}. \quad (8.13)$$

They further establish that refining the discretisation of the grid used to approximate the PDE yields convergence to the true value.

Theorem 8.2 (Global Convergence; Theorem 2.8 in [CFLSY20]). *Assume that x and y are piecewise linear*

paths as above satisfying that

$$\sup_{I \times J} |\langle \dot{x}_s, \dot{y}_t \rangle| \leq M \quad (8.14)$$

for some $M > 0$. Then there exists a constant $C_1 > 0$, depending on M and $K_{x,y}$, such that for any $\lambda \geq 0$ the numerical solution \hat{K} obtained by applying the scheme define by (8.12) on P_λ satisfies

$$\sup_{I \times J} |K_{x,y}(s, t) - \hat{K}(s, t)| \leq \frac{C_1}{2^{2\lambda}}. \quad (8.15)$$

In [CFLSY20] it was found that rescaling the paths x and y ensures that coarse partition choices of $\lambda = 0, 1$ were sufficient to provide good approximations, as illustrated in Figure 10.




Figure 10: Example of error distribution of $k_{x,y}(s, t)$ on the grids P_0 , P_1 and P_2 - Figure 2 in [CFLSY20]. Copyright ©2021 Society for Industrial and Applied Mathematics. Reprinted with permission. All rights reserved.

It is remarked in [CFLSY20] that the “PDE structure” allows for efficient GPU implementation of the scheme via parallelization. Figure 11, which is Figure 3 in [CFLSY20], illustrates the complexity reduction from quadratic on CPU to linear on GPU.

In [CFLSY20], the untruncated signature kernel is compared against other kernels via the task of classification with Support Vector Machine (SVM) classifiers [Vap98]. SVM classifiers have been successfully used in medicine [BCDFHS00], image retrieval [CT01] and text classification [KT01], for example. Given $X = \{x_1, \dots, x_n\}$ and a reproducing kernel k on X with associated reproducing kernel Hilbert space (RKHS) H_k , consider the pairs $\{(x_i, y_i) : i \in \{1, \dots, n\}\}$. For binary classification we have $y_i \in \{-1, 1\}$ and the binary SVM classification algorithm aims at solving the minimisation problem

$$\min_{f \in H_k} \sum_{i=1}^n L(y_i, f(x_i)) + \lambda \|f\|_{H_k} \quad (8.16)$$

where $L(y_i, f(x_i)) := \max\{0, 1 - y_i f(x_i)\}$ and λ is a penalty hyperparameter. The optimal solution f_* is given by

$$f_*(x) := \text{sign} \left(\alpha_0 + \sum_{i=1}^n \alpha_i y_i k(x, x_i) \right) \quad (8.17)$$

where $\alpha_0, \dots, \alpha_n$ are scalar coefficients solving a certain quadratic programming problem, see [SS18] for full details.

Selecting an appropriate kernel when X consists of multivariate time-series is difficult [SS09]. When all time-series in X have the same length, the standard kernels (i.e. linear, polynomial, Gaussian etc) work well. But outside this setting, kernels specifically designed for sequential data are required. The performances of an




Figure 11: Comparison of the elapsed time(s) to reach an accuracy of 10^{-3} from a target value obtained by solving the signature kernel PDE on a fine discretization grid (P_5). Each run simulates 5 (piecewise linear interpolations of) Brownian paths. The left plot shows the dependency on the length of two paths in dimension 2. The right plot shows the dependency on the dimension of two paths of a fixed length. - Figure 3 in [CFLSY20]. Copyright ©2021 Society for Industrial and Applied Mathematics. Reprinted with permission. All rights reserved.

SVM classifier equipped with a range of different kernels on various multivariate time-series UEA datasets [BBKLL17] are presented in Table 7.

Datasets/Kernel	Linear	RBF	GAK	Sig(n)	Sig-PDE
ArticularyWordRecognition	98.0	98.0	98.0	92.3	98.3
BasicMotions	87.5	97.5	97.5	97.5	100.0
Cricket	91.7	91.7	97.2	86.1	97.2
ERing	92.2	92.2	93.7	84.1	93.3
Libras	73.9	77.2	79.0	81.7	81.7
NATOPS	90.0	92.2	90.6	88.3	93.3
RacketSports	76.9	78.3	84.2	80.2	84.9
FingerMovements	57.0	60.0	61.0	51.0	58.0
Heartbeat	70.2	73.2	70.2	72.2	73.6
SelfRegulationSCP1	86.7	87.3	92.4	75.4	88.7
UWaveGestureLibrary	80.0	87.5	87.5	83.4	87.0

Table 7: Test set classification accuracy (in %) on UEA multivariate time series datasets - Table 1 in [CFLSY20]

The *global alignment kernel* (GAK) [Cut11] is the only other kernel proposed for sequential data, and depends on a hyperparameter $\beta \in (0, 1]$. The full untruncated signature kernel (Sig-PDE) SVM is among the top 2 classifiers across all but two of the datasets. Sig-PDE systematically outperforms its truncated counterpart Sig(n). Further numerical applications of the signature kernel can be found in section 5 of [CFLSY20].

Subsequent to the appearance of [CFLSY20] *General Signature Kernels* (GSKs) have been introduced in [CLX21]. Recall from (2.26) that for given $A, B \in T(V)$ we have that $\langle A, B \rangle_{T(V)} = \sum_{n=0}^{\infty} \langle \pi_n(A), \pi_n(B) \rangle$. Loosely, the GSKs proposed in [CLX21] are obtained by replacing the inner product $\langle \cdot, \cdot \rangle_{T(V)}$ by $\langle A, B \rangle_{\phi} := \sum_{n=0}^{\infty} \phi(n) \langle \pi_n(A), \pi_n(B) \rangle$ for coefficients $\phi(n) \in \mathbb{C}$ that are *not* required to be 1. Of course $\langle \cdot, \cdot \rangle_{\phi}$ is not necessarily an inner product on $T(V)$, but one may still define an associated GSK by

$$K_{x,y}^{\phi}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{\phi} = \sum_{n=0}^{\infty} \phi(n) \langle S_{a,s}^n(x), S_{c,t}^n(y) \rangle_{V^{\otimes n}} \quad (8.18)$$

for $x \in \mathcal{V}_1([a, b]; V)$, $y \in \mathcal{V}_1([c, d]; V)$, $s \in [a, b]$ and $t \in [c, d]$. In many situations GSKs can be interpreted as an average of PDE solutions, and may consequently be computed via suitable quadrature rules [CLX21]. The extension of this analysis allows the authors of [CLX21] to obtain closed-form formulae for expressions involving the expected (Stratonovich) signature of Brownian motion and to articulate a connection between signature kernels and the hyperbolic development map.

9. Neural CDEs and Neural RDEs

The aim of this section is to discuss *Neural Controlled Differential Equations* (Neural CDEs) [FKLM20] and their subsequent development to *Neural Rough Differential Equations* (Neural RDEs) [FKLMS21]. Loosely speaking, both are techniques in which Neural networks are utilised to learn solutions to differential Equations. We begin with a brief discussion of *Neural Ordinary Differential Equations* (Neural ODEs).

Neural ODEs are a popular and successful ODE based approach to modelling temporal dynamics. They aim to approximate a mapping $F : \mathbb{R}^q \rightarrow \mathbb{R}^l$ by learning linear maps $l_\theta^1 : \mathbb{R}^q \rightarrow \mathbb{R}^d$, $l_\theta^2 : \mathbb{R}^d \rightarrow \mathbb{R}^l$, and a function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$, all depending on parameter θ , for which the following is true. Given $x \in \mathbb{R}^q$, let $z : [0, T] \rightarrow \mathbb{R}^d$ denote the solution to the ODE

$$z(0) = l_\theta^1(x) \quad \text{and} \quad z(t) = z(0) + \int_0^t f_\theta(z(s)) ds \quad (9.1)$$

for every $t \in (0, T]$. Then the value $y := l_\theta^2(z(T)) \in \mathbb{R}^l$ gives a good approximation for $F(x) \in \mathbb{R}^l$, i.e. $y \approx F(x)$. Explicit dependence of f_θ on s is only obtained by including s as an extra dimension in z_s ; see Appendix B in [BCDR18]. The use of ODEs leads to the following fundamental issue. Once the parameter θ has been learnt, the solution to the ODE is determined by the initial condition. There is no direct mechanism for adjusting the trajectory to incorporate data that arrives later [FKLM20].

In [FKLM20] the authors resolve this issue via the use of CDEs, resulting in their development of Neural CDEs. We detail their approach for a fully-observed yet potentially irregularly sampled time-series $\mathbf{x} = ((t_0, x_0), (t_1, x_1), \dots, (t_n, x_n))$ for timestamps $t_0 < t_1 < \dots < t_n$ and observations $x_0, x_1, \dots, x_n \in \mathbb{R}^p$. The natural cubic spline $X : [t_0, t_n] \rightarrow \mathbb{R}^{p+1}$ with knots at t_0, \dots, t_n such that $X_{t_i} = (x_i, t_i)$ is chosen to approximate the underlying process observed via \mathbf{x} . The resulting differentiability of the control path X allows one to treat the dX_s term as $\frac{dX_s}{ds} ds$. It is justified in Appendix A of [FKLM20] that the choice of natural cubic spline for X has the minimum regularity required.

Let $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times (p+1)}$ be any neural network model depending on parameters θ and having a hidden state of size d . Let $\zeta_\theta : \mathbb{R}^{p+1} \rightarrow \mathbb{R}^d$ be any neural network depending on parameters θ . Then the Neural CDE model is defined to be the solution of the CDE

$$z_{t_0} = \zeta_\theta(x_0, t_0) \quad \text{and} \quad z_t = z_{t_0} + \int_{t_0}^t f_\theta(z_s) dX_s \quad (9.2)$$

for all $t \in (t_0, t_n]$. The integral in (9.2) is the Riemann-Stieltjes integral. This initial condition is used to avoid translational invariance. The model output can be taken as either the evolving process z or its terminal value z_{t_n} , and the final prediction should typically be given by a linear map applied to this output.

The essential difference to Neural ODEs is that (9.2) is driven by the data process X whilst the equation for Neural ODEs is driven only by the identity map $\mathbb{R} \rightarrow \mathbb{R}$. The Neural CDE is naturally adapting to incoming data; changes in X change the local dynamics of the system. This difference is illustrated in Figure 12, which is Figure 1 in [FKLM20].




Figure 12: **Left:** The hidden state is modified at each observation, and potentially continuously evolved between observations. **Right:** The hidden state in the Neural CDE model has continuous dependence on the observed data. Figure 1 in [FKLM20]. Reproduced with permission. Copyright held by authors of [FKLM20].

In Appendix B in [FKLM20] the authors establish that Neural CDEs are universal approximators. The precise statement may be found as Theorem B.14 in [FKLM20]; informally it guarantees that a continuous function from the space of time-series can be arbitrarily well-approximated locally by a linear map applied to the terminal value of a Neural CDE. The essential idea is that CDEs can be used to approximate bases of continuous functions on path space. This is itself a consequence of the *universal nonlinearity* property of the path signature (cf. Theorem 3.4 in this article) and that the path signature solves a CDE (cf. (4.2)). Full details may be found in Appendix B of [FKLM20].

The efficacy of Neural CDEs is demonstrated in [FKLM20] via the task of written character classification. The CharacterTrajectories dataset from the UEA time series classification archive [BBDFKLLS18] is considered. This is a dataset of 2858 time series, each of length 182, consisting of the x and y coordinate positions and the pen tip force whilst a Latin alphabet character is written in a single stroke. The task is to classify which of the 20 different characters are written.

Three experiments are run in [FKLM20] in which 30%, 50% and 70% of the observations are selected to be dropped. The selections are made uniformly at random and independently for each time series. This results in a dataset of irregularly sampled by completely observed time series. The Neural CDE model is compared with several ODE and RNN based models, details of which may be found in [FKLM20]. The results are summarised in Table 8 below, which is a variant of Table 1 in [FKLM20].

Model	30% dropped	50% dropped	70% dropped
GRU-ODE	92.6 ± 1.6	86.7 ± 3.9	89.9 ± 3.7
GRU- Δt	93.6 ± 2.0	91.3 ± 2.1	90.4 ± 0.8
GRU-D	94.2 ± 2.1	90.2 ± 4.8	91.9 ± 1.7
ODE-RNN	95.4 ± 0.6	96.0 ± 0.3	95.3 ± 0.6
Neural CDE	98.7 ± 0.8	98.8 ± 0.2	98.6 ± 0.4

Table 8: Test accuracy (mean% \pm std%, computed across five runs) on CharacterTrajectories - Table 1 in [FKLM20]

The randomly removed data is the same for every model and every repeat. The Neural CDE outperforms every other model considered. Moreover, the Neural CDE performance remains roughly constant as the percentage of dropped data increases, whereas the accuracy of the other models start to decrease.

The performance of Neural CDEs is investigated through further numerical experiments in [FKLM20]. In addition to performance, the ability to utilise memory-efficient adjoint-based backpropagation in Neural CDEs ensures that they use significantly less memory than other approaches [FKLM20]. It is also explained in [FKLM20] how the Neural CDE approach can be applied to partially observed data. Consequently Neural CDEs overcome some of the issues that can arise in the use of *Recurrent Neural Networks* (RNNs) for irregularly sampled or partially observed time series [FKLM20]. The subsequent work [KLMY21] develops Neural CDEs for *online use*; that is, to learn and predict in real-time where new data arrives during inference. The key component is to replace the use of a natural cubic spline for the control path X by a new control signal satisfying four ideal requirements; see [KLMY21] for full details.

The Neural CDE approach developed in [FKLM20] numerically solves the resulting CDE via evaluations of the control path X . Consequently, Neural CDEs performance drops for long time series, and the large number of forward operations within each training epoch result in prohibitive training time [FKLMS21]. In [FKLMS21] these issues are overcome by using the log-ODE method (cf. Section 4 in this article) to numerically solve the resulting CDE. The main idea is that the log-ODE method offers a way to update the hidden state of a Neural CDE over large intervals, dramatically reducing the effective length of the time series [FKLMS21]. An additional benefit of this approach is its removal of the requirement for the control path X to be differentiable. The resulting approach is termed *Neural Rough Differential Equations* (Neural RDEs). Neural RDEs still make use of the memory-efficient continuous-time adjoint backpropagation; the resulting decrease in memory requirement is increasingly relevant as the time series length increases.

We detail the Neural RDE approach for a fully-observed yet potentially irregularly sampled time-series $\mathbf{x} = ((t_0, x_0), (t_1, x_1), \dots, (t_n, x_n))$ for timestamps $t_0 < t_1 < \dots < t_n$ and observations $x_0, x_1, \dots, x_n \in \mathbb{R}^p$. Construct a piecewise linear interpolation $X : [t_0, t_n] \rightarrow \mathbb{R}^{p+1}$ such that, for each $i \in \{0, \dots, n\}$, we have $X_{t_i} = (t_i, x_i)$.

For a choice of positive integer m much smaller than n , pick points $t_0 = r_0 < r_1 < \dots < r_m = t_n$. Whilst it is not a requirement, it is typical to choose these points to be spaced equally apart. The choice and spacing are hyperparameters of the model. For a chosen depth hyperparameter $N \in \mathbb{Z}_{\geq 1}$ we compute the truncated log signature $\text{LogSig}_{r_i, r_{i+1}}^{(N)}(X)$ for each $i \in \{0, \dots, m-1\}$. The log-ODE method (see Section 4 in this article) is then used to numerically solve the CDE (9.2) on the interval $[t_0, t_n]$ by inductively solving it over the interval $[r_i, r_{i+1}]$. That is, let $i \in \{0, \dots, m-1\}$ and assume that the solution z_t has already been defined for $t \in [0, r_i]$. Consider the equation specified in (9.2) over the interval $[r_i, r_{i+1}]$ with initial condition z_{r_i} . Then the solution is extended to interval $[t_0, r_{i+1}]$ by defining $z_\tau := \text{LogODE}(z_{r_i}, f_\theta, \text{LogSig}_{r_i, r_{i+1}}^{(N)}(X), \tau)$ (cf. (4.7)) for $\tau \in (r_i, r_{i+1}]$. Recall from Section 4 that this only involves solving an ODE over $[r_i, r_{i+1}]$, which may be done using standard ODE solvers. Figure 13 below gives a high level comparison of the CDE and RDE based approaches; it appears as Figure 1 in [FKLMS21]

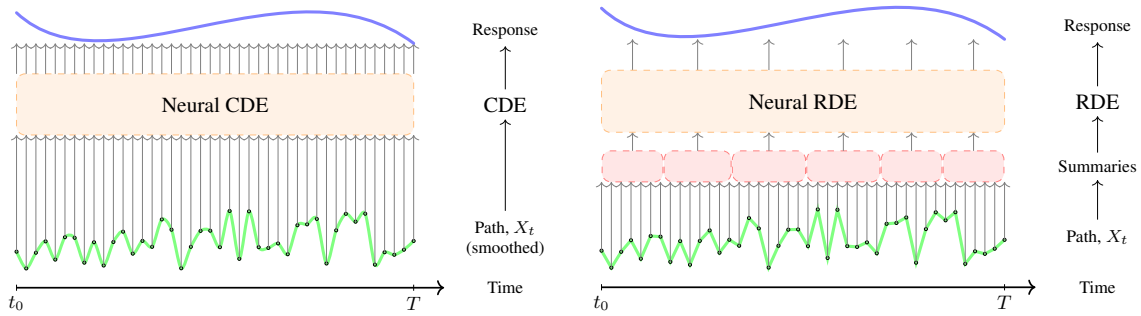


Figure 13: **Left:** CDE approach proposed in [FKLM20] in which data is smoothly interpolated and pointwise derivative information is used to drive the CDE **Right:** RDE approach proposed in [FKLMS21] in which local interval summarisations of the data given by truncated log signatures are computed and used to drive the response over the interval. Figure 1 in [FKLMS21]. Reproduced with permission. Copyright held by authors of [FKLMS21].

Numerical experiments are conducted in [FKLMS21] to illustrate the benefits of the Neural RDE approach in the setting of long time series. One task considers the EigenWorm dataset from the UEA archive [BBKLL17]. This consists of time series of length 17 984 and 6 channels (including time) corresponding to the movement of a roundworm. The task is to classify each worm as either wild-type or one of four mutant-type classes.

The regular sampling allows the choice of $t_i := i$. Neural RDEs are run for all depths $N \in \{2, 3\}$ and all step sizes in $\{2^j : j \in \{1, \dots, 10\}\}$. A Neural CDE model and an ODE-RNN model (introduced in [CDR19]) are used for baseline comparison. Precise details of hyperparameter selection, optimisers, and other model parameter choices may be found in Appendix C of [FKLMS21].

The experiment is run three times for each model and each hyperparameter combination, with the mean

and standard deviation of the accuracy recorded. The results are presented in Table 9 below, which is a version of Table 1 in [FKLMS21].

Model	Step	Test Accuracy (%)	Memory (Mb)
ODE-RNN	4	35.0 ± 1.5	3629.3
	32	32.5 ± 1.5	532.2
	128	47.9 ± 5.3	200.8
Neural CDE	4	66.7 ± 11.8	46.6
	32	64.1 ± 14.3	8.0
	128	48.7 ± 2.6	3.9
Neural RDE (depth 2)	4	83.8 ± 3.0	180.0
	32	67.5 ± 12.1	28.1
	128	76.1 ± 5.9	7.8
Neural RDE (depth 3)	4	76.9 ± 9.2	856.8
	32	75.2 ± 3.0	134.7
	128	68.4 ± 8.2	53.3

Table 9: Test accuracy (mean% \pm std%, computed across three repeats) and memory usage on EigenWorms - Part of Table 1 in [FKLMS21]

For brevity the results are only displayed for a subselection of step sizes considered; the full results can be found in Table 7 in Appendix D of [FKLMS21]. Amongst the results included in Table 9 we observe that the best accuracy results are returned by Neural RDE models, and that the accuracy achieved is a noticeable improvement on that resulting from either of ODE-RNN or Neural CDE models. This improvement over Neural CDE performance comes at a cost of Neural RDE models being more memory intensive. However Neural RDE models remain significantly less memory intensive than the ODE-RNN models.

In [FKLLO21] Neural CDEs are combined with *Neural Stochastic Differential Equations* (Neural SDEs) (see [RT19, CDLW20, HHMR20], for example) to show that the classical approach to fitting SDEs can be approached as a special case of (Wasserstein) GANs. This direct extension makes no reference to either prespecified statistics or density functions. As a consequence, in the infinite data limit *any* SDE may be learnt. Full details of this approach and several subsequent applications may be found in [FKLLO21]. Moreover, the follow-up work [FKLL21] introduces several technical innovations to improve both model performance and training speed for Neural SDEs.

The Neural CDE notebook contains an introduction to the technology developed in both [FKLM20] and [FKLMS21]. In particular, the use of the *torchcde* package for time series classification and use of the log-ode method are demonstrated. By working through the provided examples the reader will hopefully gain the experience and familiarity required to utilise the Neural CDE approach within their own work/projects.

10. Speech Emotion Recognition

Recognising emotions from audio streams has numerous real-world applications, ranging from the use of voice-based assistants such as Alexa, Siri and Google Home, to being used to assist the detection of psychiatric disorders such as bipolar disorder. Speech emotion recognition (SER) can use both voice characteristics and linguistic content to extract emotion.

A key step in SER is determining an effective and efficient representation for emotional utterances or speech segments. The complexity of emotional expressions make this challenging; the same phrase can indicate completely different emotions depending on the tone.

Many SER systems extract frame-level acoustic features (e.g. frequency, zero crossing rate, jitter, etc), called Low-level Descriptors (LLD) for utterances of various lengths, and then apply a set of statistical pooling functions (mean, max, variance, etc) to obtain fixed-size utterance-level features. Whilst the global characteristics are captured via such pooling functions, temporal variations of speech signals are not effectively extracted, diluting important regional information [AP17].

Recently, various types of deep learning models have been used to effectively model such temporal information, including both convolutional and recurrent neural networks (CNN and RNN respectively) [NV17, BMZ17, LT18, CHLRSW18]. The complex network architectures and large number of parameters involved mean these models are difficult to build and tune, time-consuming to train, and often require expensive computing resources.

In [LLNNSW19], the authors focus on the acoustic characteristics of the speech signal in order to recognise underlying emotions. Motivated by the inherent sequential structure of emotions conveyed by speech, the authors explore the use of path signatures for modelling temporal sequences of emotional utterances. They demonstrate that this method incorporates both the short-term characterisation at the frame-level and the long-term aggregation at the utterance-level. Their approach operates on minimally hand-engineered filter-bank energy features to help avoid the overfitting issues encountered by other alternative approaches, see [BMZ17].

Inspired by [JLNY16], a hierarchical tree structure is used for path signature features and tree-based convolutions are adopted for both the integration of global, regional and local information, and for filtering irrelevant and redundant information. The use of dyadic intervals allows a sufficiently fine description of both the global and local information to be captured, whilst avoiding the exponential growth in dimensionality associated with higher order signature terms.

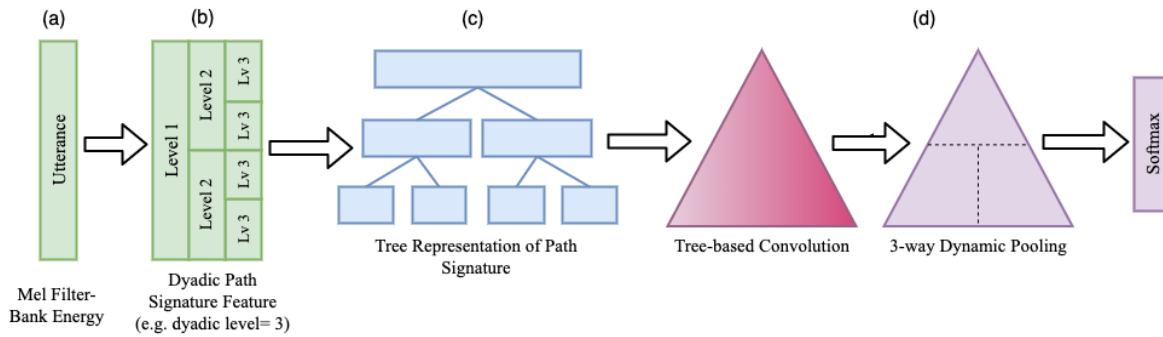


Figure 14: Overview of the proposed SER system - Figure 1 in [LLNNSW19]. Reprinted with permission.

Figure 14 illustrates an overview of the proposed SER system as explained below.

- (a) A stream of frame-level energy features is extracted from an utterance
- (b) The entire stream of frames within the utterance is segmented to dyadic paths. A truncated signature is extracted from each dyadic path
- (c) The collection of the signatures of each dyadic path is transformed to the dyadic path-tree signature representation
- (d) Tree-based convolution and dynamic pooling are applied to learn the underlying structure. An output layer is added for final classification.

The signatures are truncated to order 3, and a dyadic level of 4 is selected. The collection of all dyadic pieces forms the *dyadic path-tree* from which features at different resolutions and structural information is more easily extracted. Each node is given an associated weight matrix depending on the nodes relative position, before dynamic pooling is used to pool the features, again depending on their position in the tree. A final softmax layer is added for classification, see [LLNNSW19] for full details.

The four emotional categories *Angry*, *Happy*, *Neutral* and *Sad* form the classes for the classification. Effectiveness is evaluated using data from the Interactive Emotional Dyadic Motion Capture (IEMOCAP) database [BBCKLLMN08]. It comprises approximately 12 hours of audio-visual recordings performed by 10 actors. Each recording is split in to 5 sessions, and each session is composed of two actors, one male and one female. Overall it contains 10039 (manually segmented) utterances with an average duration of 4.5 seconds. The database can be further divided into an improvised speech data set and a scripted data set.

In [LLNSW19], the improvised data set is chosen since the scripted speech exhibits strong correlation with the manually labelled emotions leading to bias over linguistic content learning.

The openSMILE toolkit [ESW10] is used for extracting 40-dimensional features from each utterance, with an additional dimension added for time. A *leave-one-speaker-out* scheme is chosen for evaluation. This means that for each instance A, training is carried out using all instances except A, before the learnt model is then evaluated on A itself. The cross-entropy cost is considered for training, and the precise hyperparameters used are listed in table 1 of [LLNSW19]. The final SER performance is evaluated using widely adopted metrics: weighted accuracy (WA), which is the overall classification accuracy; and unweighted accuracy (UA), which averages accuracy of each emotion category. A range of neural network models are considered as baselines, and the results are summarised in the Table 10.

Model	UA	WA
COVAREP [GLMS16]	51.84	49.64
DNN-ELM [LT15]	52.13	57.91
LSTM (Speech) [GLMS16]	51.85	51.94
LSTM (Glottal) [GLMS16]	54.56	52.82
Attentive CNN [NV17]	56.83	61.95
PTS-CNN [LLNSW19]	53.03	58.90

Table 10: Model Performances - Table 2 in [LLNSW19]

The relatively simple Path-Tree-Signature based CNN (PTS-CNN) model in [LLNSW19] achieves comparable results to networks of complex design. The PTS-CNN model can deal with utterances of variable length without preprocessing and uses the openSMILE toolkit [ESW10] to obtain its features. This is in stark contrast to the Attentive-CNN model, which requires utterances to be processed to all have the same length and uses the tailor-made eGeMAPS feature set [ABDEELNSSST16]. With minimal model tuning, or manual engineering, the PTS-CNN model yields comparable results to complex neural network models making use of a range of heavily engineered emotion features (COVAREP, LSTM [GLMS16]).

11. Health Applications

11.1. Bipolar and Borderline Personality Disorders

Historically, the diagnosis of psychiatric disorders has been hampered by the inherent inaccuracy of retrospective recall of mood states. The influx of development in mobile technology (such as mobile phones, smart watches, fitbits etc) has allowed momentary assessment to obtain more precise measures of psychopathology and highlighted the shortcomings of current diagnostic categories [AGGLS18]. The NIMH Research Domain Criteria (RDoC) propose a new data-driven approach bottom up approach to diagnosis. The oscillatory nature of psychiatric symptomatology poses a significant analytic challenge.

In [AGGLS18], signature-based machine learning models are proposed to analyse data obtained from a clinical study in [BCDOPSTV16]. The study explored daily reporting of mood in participants with bipolar disorder, borderline personality disorder and healthy volunteers. The two problems tackled were the classification problem of classifying participants on the basis of their mood, and the time-series forecasting problems of predicting the participants mood the following day.

Mood data was captured from 130 individuals; 48 of whom were diagnosed with bipolar disorder (BD), 31 were diagnosed with borderline personality disorder (BPD) and 51 were healthy. For a minimum of two months, the participants rated their mood across six different categories (anxiety, elation, sadness, anger, irritability and energy) using a 7-point Likert scale with values from 1 (not at all) to 7 (very much). The data was split into streams of 20 consecutive observations, which did not necessarily arise from 20 consecutive days due to failure of participants recording. This generated 733 streams of data, which were randomly split into 513 training instance streams and 220 testing instance streams.

The streams were identified, after normalisation, as 7-dimensional paths (one dimension for time, 6 dimensions for the moods) and giving the input-output pairs $\{(R_i, Y_i)\}$ where R_i is the 7-dimensional path of

participant i . Figure 15 shows the normalised anxiety scores of a participant with BD and the associated two-dimensional path. Intuitively, if this two-dimensional path has an upward trend then the participant scores are greater than 4 in that category. Conversely, a downward trend corresponds to scores lower than 4. Drastic mood changes from day to day will result in a highly oscillating normalised path, whereas a more stable day to day mood will result in a more stable normalised path.




Figure 15: Normalised anxiety scores of a participant with bipolar disorder (above), which were calculated using the reported scores (below). High levels of reported scores correspond to upward trends, low levels of reported scores correspond to downward trends and periods of time of high oscillations in the reported scores are represented by oscillations in the path - Figure 1 in [AGGLS18]. Unaltered. Licensed under a Creative Commons Attribution 4.0 International License.

Figure 16 shows the normalised scores of each category plotted against all other categories, clearly illustrating the order in which a participants mood changes. For example, consider the Angry vs. Elated plot (third row, first column). The path starts at $(0,0)$ and initially moves left, showing that the participant is becoming less elated while their anger remains almost constant. Suddenly the period of low elation stops and the participant begins recording lower anger levels, and these low levels of anger persist for the rest of the path. On the other hand, the Angry vs. Irritable plot (fifth row, first column) is essentially a straight line, revealing that the levels of anger and irritability are highly correlated and roughly equal for this participant.




Figure 16: Normalised scores of each category plotted against all other categories, for a participant with bipolar disorder. The red point indicates the starting point and each plot has its own scale - Figure 2 in [AGGLS18]. Unaltered. Licensed under a Creative Commons Attribution 4.0 International License.

The truncated signatures $\Pi_n(S(R_i)) = S^{(n)}(R_i)$ for $n = 2, 3, 4$ are used to provide new input-output pairs. These truncated signatures will, in particular, capture the order of a participants mood change as illustrated in Figure 16. The strategies for each task are broadly similar.

For the classification, the participant of interest was removed from the data before softmax regression was used on the remaining pairs $\{(S^{(n)}(R_i), Y_i)\}$. Removing the participant of interest from the training reduced the risk of overfitting and increased the models robustness to new data. The model was then tested on 20-observations periods from the participant of interest, with the proportion of the periods of time for which they were classified as each of healthy, bipolar or borderline personality recorded.

Earlier work [BCDOPSTV16] already established there are differences in mean mood scores between the overall groups, hence classifying the streams of 20 consecutive observations on the basis of comparison with the mean score in each mood category was used as a baseline model for comparison.

The signature method significantly out-performed the naive mean-based model; the mean-based model classified 54% of participants correctly, whilst the signature model correctly classified 74.85% of participants. Incorrect predictions were mainly found in the split between bipolar disorder and border personality disorder participants, whereas the model very clearly distinguished healthy participants from the group. Full details may be found in [AGGLS18].

For the mood prediction, mean absolute error (MAE) regression was used on the input-output pairs $\{(S^{(n)}(R_i), Y_i)\}$. For testing, a prediction \hat{Y}_i was deemed correct if $|\hat{Y}_i - Y_i| \leq 1$, reflecting that we are only interested in predicting the correct class, rather than capturing the exact score. The model is benchmarked against the simple model of predicting that the next days score will be the same as the previous. The results are summarised in the following table. More detailed performance results for the signature-based method can be found in Table 3 in [AGGLS18]

Model	Healthy Accuracy (%)	Bipolar Accuracy (%)	Borderline Accuracy (%)
Same-as-prior	61-92	46-67	44-62
Signature	89-98	82-90	70-78

Table 11: Comparison of mood prediction models - Summary of Table 3 in [AGGLS18]

The subsequent work [LLNSTWW20] extends the use of path signature techniques to the classification of Bipolar disorder and Borderline Personality disorder using multi-modal datasets. The dataset considered comes from the Automated Monitoring of Symptoms Severity (AMoSS) study [AGGLS18, BCDOPSTV16] in which volunteers self-monitored their own daily mood via a smartphone app in conjunction with a range of wearable devices. Among the 139 participants, 53 had been diagnosed with BD, 33 had been diagnosed with BPD, and 53 had neither (termed healthy control, HC). Halfway through the study, 62 participants were interviewed by 2 clinicians and 2 psychology graduates to gather feedback about the scheme. The interviews were either conducted in-person or by telephone. The interviews were semi-structured with the topics remaining within the scope of

- Users experience of the smartphone app,
- Users experience of the wearable device,
- Benefits of participation in the study, and
- Potential improvements.

The topics were non-clinical, with no direct effort made to establish either a diagnosis or the participants mental state at the time of the interview.

The AMoSS interview (AMoSS-I) dataset considered in [LLNSTWW20] consists of 50 randomly sampled interviews that were initially transcribed by the same interviewers. An alignment procedure was applied to the audio recordings and manual transcribed text, as detailed in [LLNSTWW20]. In total, 67 features are extracted from the processed dataset including 28 linguistic features denoted as **LING** (such as mean sentence length, number of first person pronouns), 19 semantic content features **CNT** (such as number of words categorised as a certain emotion), 6 dialogue features **DIAL** (such as number of short pauses of less than half a second, number of pauses of longer than half a second), 3 features dealing with interruptions and talking over others, and 11 averaged features representing each pause. Full details of these features their extraction may be found in Section 3 of [LLNSTWW20].

A *leave-one-participant-out* evaluation scheme is selected and combined with logistic regression with the aim of classifying whether a participant was BD, BPD or HC based on their interview. The signature transform truncated to order 3 was applied to the path given by each extracted feature. The correlation of the features in the training data were ranked via a psychological scale (so-called International Personality Disorder Examination (IPDE) score, see [ABBCCDDFLS94]) and only those features with a p -value below 0.001 were retained. The selected features are then fed to the classifier for 3 separate binary tasks:

- BD vs. HC,
- BPD vs. HC, and
- BD vs. BPD.

Three separate experiments are conducted; one in which features are only extracted from the speech of the participant, one in which features are only extracted from the interviewer, and one in which features are extracted from the speech of both.

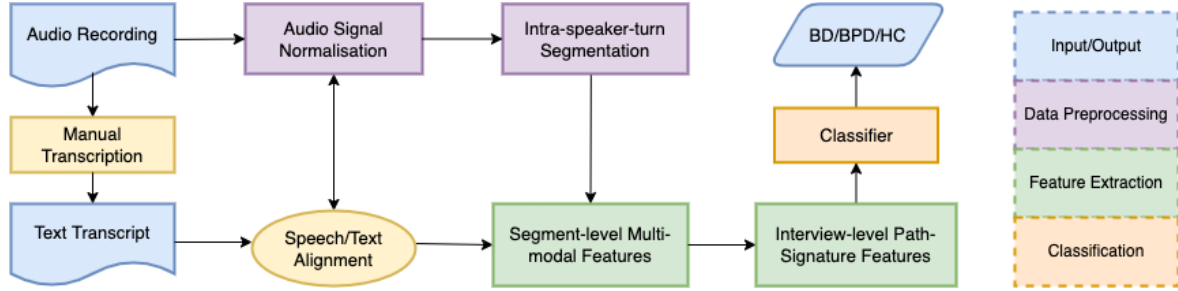


Figure 17: Data pipeline including three stages of data preprocessing (in purple), feature extraction (in green) and classification (in orange) - Figure 2 in [LLNSTWW20]. Reprinted with permission.

The results for the classification task are summarised in Table 12.

Subject	AUROC		
	BD vs. HC	BPD vs. HC	BD vs. BPD
Participant	0.810	0.733	0.817
Interviewer	0.304	0.473	0.231
Both	0.494	0.431	0.657

Table 12: Average AUROC for each binary classification - Table 3 in [LLNSTWW20]

The scores presented are the average of the AUROCs across all interviews. Extracting features from the participant alone results in a good classifier for the task. Given that the interviews were non-clinical and not aimed at establishing the mood of a participant, it is not surprising that the features extracted from the interviewer perform badly.

Ablation experiments were conducted to examine the affect of removing each feature type. The results are summarised in Table 13.

Features	BD vs. HC	BPD vs. HC	BD vs. BPD
All	0.810**	0.733**	0.817**
All - CNT	0.810**	0.733**	0.787**
All - Dial	0.768**	0.733**	0.811**
All - LING	0.625*	0.578*	0.669*
All - LING - CNT	0.642*	0.703*	0.604*
All - LING - Dial	0.442*	0.429*	0.550*
All - CNT - Dial	0.768**	0.733**	0.763**

Table 13: Feature ablation results (AUROC) for each task; p -value used for feature selection: ** < 0.001 and * < 0.005 - Table 4 in [LLNSTWW20]

The **LING** features are the biggest contributors; in fact, when they are removed the p -value threshold has to be made 5 times larger to allow *any* features to be selected. The results illustrate that combining different types of features works better than relying on any individual type.

The subsequent work [LLSVWW21] investigates the use of the invisibility reset transform (cf. (2.15)) to effectively model the conversation dynamics by enhancing the signature representation of conversation speech.

11.2. Alzheimer's Disease

Alzheimer's disease (AD) is the most common form of dementia in older people, affecting 6% of the population aged over 65. Memory loss is a mild cognitive impairment (MCI) commonly indicating early stage of AD. MCI is diagnosed with no evidence of dementia, and MCI does not always progress to dementia. However, the progressive and irreversible loss of brain function caused by AD suggests that potential drug therapies should be used as early as possible. Hence there is a demand for reliable early predictions of which individuals will develop AD.

Using brain imaging data to derive features for predicting a diagnosis of AD has been investigated in many works, see [ABDGKKMMRRS15, ABILNPS17, KLMMSSSTVWWW11] for example. In [GLM19], signature-based techniques are developed to distinguish between healthy and Alzheimer groups from preprocessed MRI data. Data from the TADPOLE grand challenge competition in 2019 was used for the classification tasks considered in [GLM19]. The 1737 participants from the TADPOLE data set are split into three groups; the 688 who have a diagnosis of Alzheimer's disease at some time (labelled *AD*), the 424 who always have a healthy diagnosis (labelled *NL*), and the 484 who always have an MCI diagnosis (labelled *MCI*). This results in the comparison of participants whose diagnosis converts to Alzheimer's disease with those whose diagnosis remains unchanged.

The training data is selected in [GLM19] by selecting participants with a first diagnosis of AD at 36 months from baseline (start of monitoring). Each selected participant is further required to have at least four measurements of all the variables WholeBrain (*w*), Hippocampus (*h*) and Ventricles (*v*) (see [GLM19] for details of these variables) in the 24 months since baseline, with one measurement at the 24 month time point. Finally, each selected participant from the *AD* set is required to have matching counterparts in both the *NL* and *MCI* sets. Given a member $x \in AD$, an individual $y \in NL$ or $y \in MCI$ qualifies as a counterpart x if the following requirements are satisfied. Firstly, the ages of y and x must differ by no more than 5 years. Secondly, the diagnosis for y must have remained unchanged for the 72 months since their first (baseline) measurement. Finally, y must have at least four measurements of all the variables *w*, *h*, and *v* up to month 24, again including a measurement at month 24 itself.

The test data is selected making use of measurements starting from 12 months since the first baseline measurement. To reflect the real world the test data imposes no matching/counterparting between individuals in *AD* and those in *NL/MCI*. The measurements used for analysis of the test data are at 12, 24 and 36 months for all three sets (*AD*, *NL*, *MCI*).

Two classification tasks are considered in [GLM19]; *AD* vs. *NL* and *AD* vs. *MCI*. Binary logistic regression, which models the log probabilities of the outputs as linear functions of the inputs. This results in the model being more easily interpreted than more sophisticated classifiers such as random forests. The input features are selected via Lasso regularisation. The input for LASSO is a vector formed from the three variables, WholeBrain (*w*), Hippocampus (*h*) and Ventricles (*v*), and either the signature or the log-signature of the path determined by the variables *v*, *h*, and *w*. The training in [GLM19] uses 10-fold cross-validation. The LASSO regularisation coefficient λ is increased to result in a sparse set of variables to act as predictors; see [GLM19] for full details.

The features selected for each classification task are summarised in the following table which is a version of Table 2 in [GLM19].

Task	Signature Feature Set	Log Signature Feature Set
AD vs. NL	Hippocampus-BL [Incr.Ventricles] (Hippocampus,Time) (Hippocampus,Time) (Hippocampus,WholeBrain)	Hippocampus-BL [Incr.Hippocampus] [Incr.Ventricles]
AD vs. MCI	Hippocampus-BL Ventricles-BL (Hippocampus,Time) (Hippocampus,WholeBrain) (Time,Ventricles) (Hippocampus,Hippocampus)	Hippocampus-BL Ventricles-BL [Incr.Hippocampus] [Incr.Ventricles] (Time,Ventricles)

Table 14: Comparison of feature sets selected for each classification task in [GLM19] - Table 2 in [GLM19]

The notation used in Table 14 is as follows. The baseline value of a variable X is denoted by X -BL. The increment of a variable X are denoted $[Incr.X]$ and correspond to depth one signature terms. Area terms between variables X and Y are denoted by (X,Y) and correspond to particular combinations of depth two signature terms. The precise details may be found in [GLM19].

The resulting classifiers achieve at least 90% accuracy on both classification tasks considered on the test data, see Table 3 in [GLM19] for full details of the performance. Whilst the features provided by the signature method correspond to known AD pathology that can be extracted manually, the path signature provides a systematic way of automatically generating these features *without* requiring any prior knowledge of the pathology.

11.3. Early Sepsis Detection

Sepsis is an overaggressive autoimmune response to infection that can cause life-threatening damage to the body's organs. It was estimated that in 2017 sepsis affected 50 million people worldwide and caused 11 million deaths [HKLMNS20]. In America alone it is thought to be responsible for one in three hospital deaths [NCEZID16], and the cost of admission and patient care has been estimated to exceed \$41.5 billion [BSS20-I, BSS20-II, BSS20-III]. The time of detection is critically linked to the mortality rate of sepsis. In cases of septic shock, it is known that the risk of death increases roughly 10% for every hour of delay in antibiotic treatment [CFGKKLPRSSTWZ06]. Early detection of sepsis is evidently crucial to improving sepsis management and mortality rates.

The works [HKLMNS19, HKLMNS20] investigate the use of path signature techniques for early prediction of sepsis. Unlike prior existing machine learning algorithms for this task, the Sepsis-3 definition [DSS16] is used to denote the onset of sepsis. The development of the model in [HKLMNS19] was during the 2019 PhysioNet challenge entitled "Early Prediction of Sepsis from Clinical Data." Challenge participants were invited to submit algorithms that were trained on the same set of readily attainable ICU data and validated under a common performance metric on unseen test data. The data and performance metrics used are detailed in [CJJNRSSW19]. The algorithm proposed in [HKLMNS19] was the first placed entry in the official phase of the challenge [HKLMNS20].

The aim of this section is to present an overview of the method developed in [HKLMNS19]. In addition to [HKLMNS19], we will use the follow up work [HKLMNS20]. The article [HKLMNS20] is an extension of the Computing in Cardiology conference proceedings paper [HKLMNS19] which includes further method details, additional results, and discussion.

The PhysioNet/Computing in Cardiology Challenge 2019 data were sourced from ICU patients in three separate hospital systems (which we label Hospital A, Hospital B and Hospital C). The data resulting from Hospital A and Hospital B were split into a publicly available training set and an undisclosed test set, both to be used for model development and testing. The test set and the data from Hospital C remained private, with submitted models being scored against these unseen data from all three systems. The training set comprised of 40 336 patients with 40 features consisting of demographic, vital signs, and laboratory data recorded per patient. The data was indexed with time at 1-hour increments and predictions were to be made sequentially as

each hour in a patient’s time series using *only* the information observed up until that time (i.e. without making use of any future information). A custom utility function (detailed in [CJJNRSSW19]) was used to score the models.

Loosely speaking, the utility function is designed to ensure the following. Firstly, false-positives amongst patients that never develop sepsis are penalised, whilst correct true negatives score zero. Secondly, for patients that did develop sepsis, early prediction was penalised and false negatives were more heavily penalised. In addition correct true positives are rewarded. The penalisation of early prediction was included to promote predictions being made in the desired 6-hour window directly proceeding the on-set of sepsis time defined using the Sepsis-3 definition [DSS16].

The datasets are first augmented to include a number of additional features thought, based on literature review and “expert knowledge”, to be useful for discerning the onset of sepsis [HKLMNS20]. The following hand-crafted features derived from the data are additionally added (see Table 1 in [HKLMNS19]).

- **Shock Index:** the ratio of the heart rate to the systolic blood pressure.
- **BUN/CR:** the ratio of bilirubin to creatinine.
- **Counter:** variable for the temperature and the laboratory values that records the number of times a given variable is measured over a pre-determined look-back window. Its inclusion is designed to exploit the idea that measurement frequency provides an indication of patient health (i.e. an increasing in sampling rate may indicate physician concern about a patient).
- **Max/Min:** the maximum and minimum of the vital signs over a pre-determined look-back window.
- **PartialSOFA:** partial construction of the *Sequential Organ Failure Assessment* (SOFA) score [DSS16]. The resulting feature is termed PartialSOFA since the dataset did not include all variables that comprise the SOFA, meaning that the PartialSOFA score was calculated only on the required information available in the data. Hence the PartialSOFA score was calculated based on threshold conditions on each of the platelet count, bilirubin, mean arterial pressure (MAP), and creatinine variables.
- **SOFA-Deterioration:** Binary label; 1 if PartialSOFA has decreased by 2 in the last 24 hour window.

The PartialSOFA and SOFA-Deterioration variables are included since deterioration of the SOFA score is a requirement of the Sepsis-3 definition. The size of the look-back windows chosen for the Counter and Max/Min variables are treated as hyperparameters and optimised during training.

Once the data has been augmented to include these hand-crafted features, the signature transform is applied to the resulting time series. A sliding window approach is used; signature features are computed for each time point over a window of pre-determined look-back size. The input paths were augmented to include a time dimension, and the cumulative sum followed by the lead-lag transformation were applied as further augmentations prior to truncated signatures being computed. Both the truncation level and the look-back window length are treated as hyperparameters to be optimised during training.

The challenge data was prelabelled with the value 1 at any location of sepsis occurrence or predefined window around sepsis onset and zero otherwise. The method proposed in [HKLMNS19] creates an alternative labelling that accounts for information about the utility score to enable the classifier to place greater importance on points that lead to a larger score if predicted correctly. If $U_y(x, t)$ denotes the utility score of predicting y for patient x at time t , then the *modified utility score* (MUS) is defined as $U_M(x, t) := U_1(x, t) - U_0(x, t)$. It is against this labelling that the regressor is trained.

Stratified five-fold cross-validation, with a uniform distribution of time points and sepsis labels in each fold, is used for hyperparameter optimisation. Precise details may be found in [HKLMNS20]. The final values of the parameters can be found in Table 2 in [HKLMNS19]. They may be summarised as follows. The Counter variables are computed over a look-back window of size 8, whilst the Max/Min variables are computed over a look-back window of size 6. The streams given by PartialSOFA, MAP and BUN/CR are augmented with a time dimension and a lead-lag transformation, before signatures truncated to depth 3 were computed using a look-back window of size 7. For the remaining non-stationary streams, the cumulative sum augmentation is applied followed by the lead-lag transformation, before signatures truncated to depth 3 are computed using a look-back window of size 7.

The performance of the model using the final hyperparameter values is recorded via the average utility score over the five-folds used in cross-validation. On the training data and testing data from Hospital A this score was 0.442 and 0.433 respectively, whilst on the training data and testing data from Hospital B it was 0.421 and 0.434 respectively. The similar performance on both the training and testing data suggest the model was not over-fitted when restricted to the same hospital system on which it was trained.

The data from Hospital C was not included in the public training data, and was used only for validation. The model achieved an average utility score of -0.123 on the data from Hospital C. This is significantly worse than the performance on the data from Hospital A and Hospital B, highlighting the limitations to using the method trained on one hospital system to make predictions on a different one [HKLMNS20]. The utility scores for the model using the final hyperparameters may be found in Table 1 in [HKLMNS20].

The impact of the various features on the models performance are explored in [HKLMNS20]. The cross-validated and averaged utility score predictions on the training data for models trained using different subsets of features are recorded in Table 2 in [HKLMNS20]. Four different subsets of features are considered; Time only (T), Time + Original 40 features (TO), Time + Original 40 features + Nonsignature features (TON), and Time + Original 40 features + Nonsignature features + Signature Features (TONS).

The average utility scores achieved using each subset are 0.282 for T, 0.389 for TO, 0.422 for TON and 0.434 for TONS. These scores reveal that the time-only feature is the single most useful feature. Whilst the inclusion of the signature features does not result in a dramatic improvement (the score increases from 0.422 to 0.434), it nevertheless illustrates that the representation of the information after the signature transformation is beneficial to learning.

The model from [HKLMNS19] is designed to optimise the predefined utility function. An extension to allow its use in an in-hospital environment to provide clinically actionable information is proposed in [HKLMNS20]. At each time point, the larger the models output value the higher the risk of sepsis. When a specified operating point threshold is exceeded, the subject is designated as a “sepsis-risk patient”, indicating that closer monitoring or further tests may be warranted. This operating point threshold can be chosen to achieve the most clinically meaningful sensitivity and specificity [HKLMNS20].

Returning to the training set of 40 336 patients, in [HKLMNS20] the authors implement this for a range of operating point thresholds. The results are presented in a confusion matrix where a true negative represents that no call was made and the person did not develop sepsis, and a true positive represents that the patient being flagged and developing sepsis “at some point” after this call. The choice of operating point threshold so that 33% specificity (the proportion of correct positives that were correctly identified) is achieved results in the following confusion matrix.

	Predicted Sepsis	Predicted No-Sepsis
Actual Sepsis	1777	1150
Actual No-Sepsis	3554	33855

Table 15: Confusion matrix displaying the number of people predicted as likely to get sepsis compared with those who actually end up with sepsis with the threshold tuned to 33% specificity - data appears in Figure 2 in [HKLMNS20]

The signature-based model presented in [HKLMNS19, HKLMNS20] for the early prediction of sepsis offers a competitive approach to discerning early onset sepsis from health data streams. The Early Sepsis Detection notebook provides an introduction to implementing the methodology proposed in [HKLMNS19, HKLMNS20]. The data used to train the model in this notebook are the sequences of physiological and laboratory-observed measurements contained in the MIMIC-III dataset. This consists of electronic health records for 40000 patients in intensive care at the at the Beth Israel Deaconess Medical Center, Boston, Massachusetts, between 2001 and 2012. These data include, for example, patients’ heart rates, temperatures, and oxygen saturation levels, all recorded repeatedly over time for each patient. The task considered is to use the classifier to predict, at time t , whether a patient will go on to develop sepsis by time $t + T$, for some pre-determined $T > 0$, based on the patients measurement sequences recorded up until time t . The notebook allows one to workthrough a version of the methodology of [HKLMNS19, HKLMNS20] in a simplified setting to illustrate the use of signature methods for the early detection of sepsis.

11.4. Information Extraction from Medical Prescriptions

Medical prescription notes are a valuable source of important patient information that may not be recorded elsewhere. Manual extraction and annotation is time-consuming and, since it must be done by specialists, expensive. Natural language processing (NLP) tasks offer an approach to automating these tasks. Advances in pre-training large-scale contextualised language representations such as ELMo [CGILNPZ18] and BERT [CDLT18] have improved the performance of many NLP tasks. Several studies have fine-tuned BERT to clinical text (which differs substantially from general text) for NLP tasks [ABJMMNWW19, LPY19, RSWX19, AHR20]. Many tasks can be formulated as a classification or regression task, wherein either a simple linear layer or sequential models such as LSTM are added after the BERT encoding of the input text as task-specific prediction layer [CGGLLNPTU20].

BERT’s model architecture does not contain any recurrence or convolution, meaning additional positional encoding is required to model the order of the sequence (i.e. word order). In [BLNW20], the authors integrate the signature transform method into the Transformer model in order to naturally capture sequential ordering information in an effective manner.

The model architecture of BERT is composed of N identical layers. Each layer has a multi-head self-attention mechanism followed by a position-wise fully connected feed forward network. The attention function takes three input vectors query Q , key K and value V , and generates a weighted sum of the value V , with the weights given as the softmax of a rescaling of the dot-product QK^T (see Section 2 of [BLNW20] for the precise definition). Multi-head attention splits Q , K and V into multiple heads via linear projection, allowing the model to attend to information at different positions from different representation subspaces in parallel. Each projected head first goes through the scaled dot-product attention function, then concatenated, and finally projected once more to output the final values. Positional information is not explicitly modelled by the Transformer encoder without modifying its inputs to include a representation of absolute position [BLNW20].

The main extension to the Transformer architecture proposed in [BLNW20] is replacing this attention function with a *SigAttention* function. The Attention function itself is first taken, before the outputs dimension is reduced by an affine map. The truncated signature $S^{(N)}$ for some $N \in \mathbb{N}$ is taken to give the final output. A dropout rate (set to be 0.1 in [BLNW20]) is used to ensure the output dimension is the same as the input dimension. The SigAttention mechanism is illustrated in Figure 18 below.




Figure 18: Sig-Attention mechanism - Part of Figure 2 in [BLNW20]. Reproduced with permission. Colour Alterations. Licensed under a Creative Commons Attribution 4.0 International License.

The proposed *Additive Multi-Head Sig-Attention* model combines the information from the input sequence

with the output of Sig-Attention in different representation subspaces; see Figure 19 below which is part of Figure 2 in [BLNW20]. The model takes the embedding of an input sequence as well as the queries, keys and values as input. As shown in Figure 19, the queries, keys and values as well as the input embeddings are first linearly projected to have the required dimensions. Sig-Attention is then applied to the query, key and value vectors, whilst the signature is taken of the projected input embedding. The two output signatures are combined to result in one signature vector. Finally all the signature vectors are concatenated and passed onto the next sub-layer.

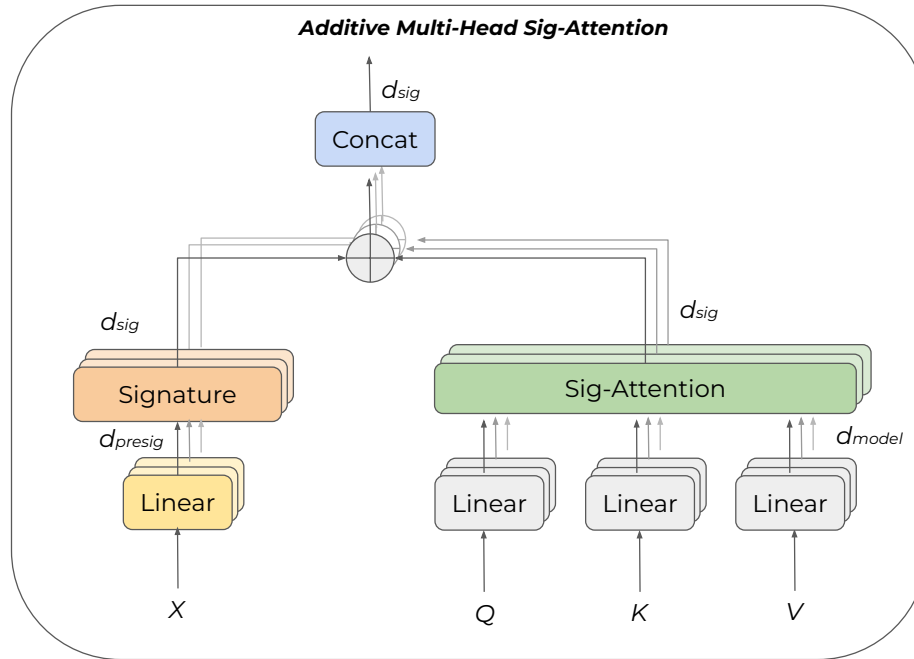


Figure 19: Additive Multi-Head Sig-Attention - Part of Figure 2 in [BLNW20]. Reproduced with permission. Colour Alterations. Licensed under a Creative Commons Attribution 4.0 International License.

A dataset of 3852 distinct prescriptions for beta-blockers, provided by Karolinska University Hospital, written in Swedish is used for experimentation. Medical practitioners have annotated three labels QUANTITY, QUANTITY TAG and INDICATION.

- QUANTITY: Total amount of capsules prescribed
- QUANTITY TAG: one of 5 classes labelling the quantity prescribed to the patient:
 - i. *Not Specified*: the quantity was not specified on the prescription
 - ii. *Complex*: a range of quantities was given and the quantity is an average
 - iii. *PRN*: prescription to take only if needed
 - iv. *As Per Previous Prescription*: refers to guidance in previous prescription
 - v. *Standard*: standard prescription
- INDICATION: one of 5 classes covering the purpose of the prescription:
 - i. *Cardiac*
 - ii. *Tremors*
 - iii. *Migraine*
 - iv. *Others*

v. *NA (Not Annotated)*

An example annotated prescription is included in Table 16. This is a reproduction of Table 1 in [BLNW20]

Swedish	Translated English	Indication	Quantity	Quantity Tag
1 TABLETT FOREBYGGANDE MOT MIGRAN	1 tablet prevention against migraine	Migraine	1	Standard
MOT HOGT BLODTRYCK OCH HJARTKLAPPNING 2 TABLETTER KLOCKGAN 08:00 1 TABLETT KLOCKAN 18:00	Against high blood pressure and heart palpitations 2 tablets at 08:00 1 tablet at 18:00	Cardiac- hypertension	3	Standard
2 TABLETTER KL. 08, 2 TABLETT KL. 20. DAGLIGEN. OBS KVALSDOSEN HAR HOJTS JFRT MED 070427	2 tablets kl. 08 2 tablet kl. 20 Daily. Note the evening box has hojts jfirt with 070427	NA	4	Standard
1 tablett vid behov mot stress	1 tablet if needed against stress	Anxiety	1	PRN
FOR BLODTRYCK OCH HJARTRYTM-	For blood pressure and heart rhythm -	Cardiac- hypertension- dysrhythmia	0	Not Specified
1 tabl pA morgonen och en halv tabl pA kvAllen fOr hjArtrytmn	1 table in the morning and a half table in the evening for heart rhythm	Cardiac- dysrhythmia	1.5	Standard
1-2 tablett 2 gAnger dagligen	1-2 tablets 2 times daily	NA	3	Complex

Table 16: Example prescriptions with translations (via Google translate) and annotations; the final three columns are the labels of interest that are sought to be extracted automatically - Reproduced variant of Table 1 in [BLNW20]

The following two approaches are proposed in [BLNW20].

- Encode the Swedish text directly using Multilingual BERT (M-BERT) [CDLT18], and
- Translate the prescriptions and then apply ClinicalBERT [AHR20] that is pretrained on clinical English.

Both methods are applied to the multi-task learning problem consisting of a regression problem (find the QUANTITY) and two classification problems (find the QUANTITY TAG and INDICATION). The proposed Sig-Transformer Encoder (STE) multi-task learning architecture is depicted in Figure 20.




Figure 20: Proposed multi-task learning architecture with Sig-Transformer Encoder (STE) - Figure 4 in [BLNW20]. Reproduced with permission. Colour Alterations. Licensed under a Creative Commons Attribution 4.0 International License.

A single linear layer is used as the classifier for QUANTITY TAG whilst two-layer networks are used for both QUANTITY and INDICATION. Cross-entropy loss is used for the two classification tasks whilst the mean-square error is used for the regression. The overall loss function is taken to be a weighted sum of the loss functions for all three tasks. The STE model is compared against both the baseline Base model (i.e. the architecture of Figure 20 with the STE step removed) and against an LSTM model given by the same architecture as Figure 20, but with the STE step replaced by an LSTM step. The results are summarised in Table 17.

Model	QUANTITY (MSE)	QUANTITY TAG (f1 score)	INDICATION (f1 score)
Base	0.50	0.60	0.08
ClinicalBERT	0.21	0.89	0.09
M-BERT	0.41	0.63	0.10
Base+LSTM	0.45	0.76	0.06
ClinicalBERT+LSTM	0.47	0.64	0.10
M-BERT+LSTM	0.50	0.68	0.02
Base+STE	0.23	0.78	0.03
ClinicalBERT+STE	0.36	0.77	0.06
M-BERT+STE	0.15	0.92	0.05

Table 17: Performance comparisons for information extraction - Table 3 in [BLNW20]

Without using translation (i.e. using M-BERT), the addition of LSTM and STE improves performance for QUANTITY and QUANTITY TAG. Overall, the M-BERT+STE model proposed in [BLNW20] performs best for QUANTITY and QUANTITY TAG. Whilst the proposed methods gives worse performance for INDICATION, the unanimous poor performance by all models on this task suggests there is an underlying issue with the class boundaries. It is likely that the aggregation of the original 44 classes to the final 5 classes has left the new classes much less distinct since each is now formed of a range of topics [BLNW20].

Further evidence that the INDICATION task is unexpectedly more challenging is the strong individual class separation achieved by the M-BERT+STE model in the QUANTITY TAG classification task, illustrated in Table 18

Model	Standard	QUANTITY TAG (f1 score)			
		APPP	PRN	Complex	NS
Base	0.71	0.50	0.10	0.76	0.95
ClinicalBERT	0.83	0.97	0.89	0.99	0.79
M-BERT	0.94	0.41	0.04	0.81	0.98
Base+LSTM	0.93	0.36	0.99	0.77	0.74
ClinicalBERT+LSTM	0.90	0.23	0.99	0.61	0.49
M-BERT+LSTM	0.29	0.99	0.72	0.47	0.93
Base+STE	0.97	0.84	0.77	0.44	0.88
ClinicalBERT+STE	0.99	0.70	0.85	0.65	0.65
M-BERT+STE	0.86	0.97	1.00	0.89	0.86

Table 18: Performance for individual classes comparison; APPP means “As Per Previous Prescription” and NS means “Not Specified” - Table 6 in [BLNW20]

The M-BERT+STE model performance is consistent across all classes, and for each class its performance is comparable to that of the best model for that particular class. In fact, every model other than ClinicalBERT+STE is considerably out-performed by M-BERT+STE on at least one class. Incorporating signature transform methods with the self-attention mechanism has enabled the authors of [BLNW20] to create a relatively simple model whose performance is at least comparable with its competitors on the information extraction tasks considered in [BLNW20].

12. Landmark-based Human Action Recognition

Human action recognition (HAR) has a wide range of applications such as video surveillance and behavioural analysis. The task is to recognise an action from video footage. *Landmark-based Human Action Recognition* regards objects as systems of correlated labelled landmarks. Recognising actions from the evolution of vectors connecting these labelled landmarks is motivated by the moving light-spots experiments in [Joh73]. They demonstrated that people can detect motion patterns and recognise actions from several bright spots distributed on a body. The challenge of LHAR is to train a computer to recognise actions based on the evolution of some set of vectors in a fixed vector space, usually \mathbb{R}^d for some $d \in \mathbb{N}$.

LHAR uses anonymous data; only the position of a certain number of markers is recorded. The use of de-identified data is often more suitable. For example, monitoring vulnerable people in their homes to spot accidents and falls. It is unlikely that many people would be happy to have full video data constantly recorded from within their homes. However, only having the location of, say, 20 markers recorded, without any person identifying features, is likely a more palatable option.

Two major challenges associated with LHAR are designing reliable discriminative features for spatial structural representation and modelling the temporal dynamics of motion. There are two main approaches to LHAR: joint-based and part-based. The joint-based approach regard the body as a set of points and attempt to capture correlation among the body joints via the pairwise distances or the joint orientations. The parts-based approach focuses on connected segments of the set of points. Most methods represent spatial poses using predefined skeletal structures, only connecting points via paths following the skeleton of the body, see [ACV14] for example.

Whilst these connections seem intuitive, they are not guaranteed to be the crucial ones for distinguishing actions. The connections discarded by imposing a skeletal structure may contain valuable non-local information. For instance, the non-local displacement between two hand points is a key feature for the action of clapping. Not using skeletal structure information ensures that any potentially vital non-skeletal connections are not discarded. A further advantage is that it makes it easier to extend the method to the actions of other objects. If we are not using skeletal structure, then we need only specify the landmark locations on the body. We can easily consider other objects by experimenting with landmark location *without* needing to understand the underlying structure of this new object.

A limitation of this framework is that the labelled data must all be within the same vector space. We need

to be able to make sense of the difference between two joints, with the distance between them often being important. This is evidently the case for locations on a body, which are naturally vectors in \mathbb{R}^3 (or \mathbb{R}^4 if time is included). But if the data consisted of recorded emotions, as was the case in Section 11.1, then we would not be in this framework; there is no sensible notion of what the difference anger - depressed should be, for example.

The evolution of the connected segments is a path, and thus using the signature to understand this path is natural. In [JLNSY17], the path signature feature (PSF) transform is used to provide a suitable feature set for spatial and temporal representation of LHAR. A pose is localised by disintegration into a collection of m -node sub-paths, with the signatures of these paths encoding non-local and non-linear geometric dependencies. This offers a resolution to the problematic trade-off between hand-designed local descriptors being insufficient to capture complex spatio-temporal dependences [BGJSZ13, FTZ13], and the deep RNN, LSTM in particular, models learning features that are not as easily interpreted.

12.1. Path Disintegration and Transformations

The disintegrations considered are *Pose Disintegration* and *Temporal Disintegration*. The pose is regarded as an ordered collection of points in \mathbb{R}^d , with the ordering chosen randomly and fixed. The pose disintegration localises the pose into all possible subposes containing m points. The inherited order converts each subpose to a unique m -node sub-path that visits each node once. These paths are not restricted to being parts of the physical body. An illustration of the pose disintegration can be seen in Figure 22. The notation PSF_m is used to refer to computing the path signature features of all possible subposes containing m points in a single frame.

The temporal disintegration splits the interval $[0, T]$ into over-lapping dyadic intervals. For $j \in \mathbb{N}$, the j^{th} dyadic level of the resulting hierarchical structure is the collection of subintervals

$$\mathcal{D}_j := \mathcal{D}_j^{\text{dyadic}} \cup \mathcal{D}_j^{\text{overlap}} \quad (12.1)$$

where

$$\mathcal{D}_j^{\text{dyadic}} := \left\{ \left[\frac{i}{2^j}T, \frac{i+1}{2^j}T \right] : i \in \{1, \dots, 2^j - 1\} \right\} \quad (12.2)$$

and

$$\mathcal{D}_j^{\text{overlap}} := \left\{ \left[\frac{2k+1}{2^{j+1}}T, \frac{2k+3}{2^{j+1}}T \right] : k \in \{1, \dots, 2^j - 2\} \right\}. \quad (12.3)$$

The inclusion of the overlapping intervals $\mathcal{D}_j^{\text{overlap}}$ prevents the interior points $\frac{i}{2^j}T$, for $i \in \{1, \dots, 2^j - 1\}$, being artificially converted to endpoints in the collection \mathcal{D}_j . This dyadic decomposition of the interval $[0, T]$ is illustrated in Figure 21.




Figure 21: On the left is the standard dyadic collections $\mathcal{D}_0^{\text{dyadic}}$, $\mathcal{D}_1^{\text{dyadic}}$ and $\mathcal{D}_2^{\text{dyadic}}$, on the right is the collections \mathcal{D}_0 , \mathcal{D}_1 and \mathcal{D}_2 . Image retrieved from https://github.com/datasig-ac-uk/signature_applications/tree/master/human_action_recognition.

Signatures are computed over all the smaller intervals, reflecting the idea that the low-order signature terms over all smaller intervals will be more informative than the higher order terms over the entire interval $[0, T]$. The choice of hierarchical level $h \in \mathbb{N}$ is a trade-off between improving efficiency and introducing local noise

over finer intervals. The notation $\mathcal{D}_j - PSF_m$ is used to refer to computing the path signature features of all possible subposes containing m over all the l^{th} level dyadic decomposition \mathcal{D}_l for $l \leq j$.

Better performance is observed for taking signatures of level two over all of \mathcal{D}_0 , \mathcal{D}_1 and \mathcal{D}_2 than taking signatures of level five over the whole time interval $[0, T]$ in [JLNSY17]. Whilst using level two truncations involves simpler and quicker computations than using level five signatures, the level two signatures on \mathcal{D}_0 , \mathcal{D}_1 and \mathcal{D}_2 also result in a smaller feature set. Assuming the path is three-dimensional, the level five signature over $[0, T]$ involves $\frac{3^6-1}{3-1} = 364$ terms. In contrast, the level two signatures over all the dyadic intervals \mathcal{D}_0 , \mathcal{D}_1 and \mathcal{D}_2 involve only $11 \times \frac{3^3-1}{3-1} = 11 \times 13 = 143$ terms, which is roughly 40% of the size.

12.2. Datasets

The proposed method is trained on four datasets; JHMDB [BGJSZ13], SBU [BCHSY12], Berkeley MHAD [BCKOV13] and NTURGB+D [LNSW16]. The JHMDB dataset consists of 928 clips, each containing between 15 and 40 frames, capturing a single person doing one of 21 actions. The 2D joint positions are manually annotated, and the loss of information due to 2D projection presents an additional challenge.

The SBU Interaction is a 3D Kinect-based dataset. Kinect sensors [Mic10] are cost effective depth cameras that are capable of providing reliable joint locations via real-time pose estimation algorithms [BCF-FKMSS13]. SBU contains a total of 282 clips categorised into 8 classes of two-actor interaction. Self-occlusion (when an object is obscured due to overlapping itself) causes measurement errors in the joint locations, which again presents an additional challenge.

The Berkeley MHAD dataset is 659 clips captured by a marker-based motion capture system. The 3D locations of 43 joints are accurately captured using LED markers. There are 384 clips performed by 7 actors that are used for training, and 275 clips performed by 5 actors that are used for testing.

The Kinect-based NTURGB+D dataset is one of the largest 3D action recognition datasets and contains 56,000 clips categorised into 60 classes. There is no constraint on the number of actors appearing in each clip which, together with the large variation in viewpoint, pose a significant challenge for analysing this dataset.

12.3. Method Implementation

Gaussian noise is added over joint coordinates to simulate the errors caused by estimation. Biometric differences are compensated for by normalising the coordinate value ranges to $[-1, 1]$ over the entire clip. Similarly, each feature is normalised to the interval $[-1, 1]$. To obtain a fixed length input, $M = 10$ clips are uniformly sampled from each clip.

The pose disintegrations for $m = 2$ and $m = 3$ are considered. Since the absolute position may be essential in some applications (static CCTV monitoring for example) the *Invisibility-reset transformation* is used to retain this information in the signature. Further, multi-delayed no-future-pause Lead-Lag transformations with no-future-pause are considered, with those with smaller delays encoding short-term dependencies, and those with longer delays encoding long-term dependencies. See Section 2.2 for definitions of these transformations and Figure 22 for an illustration of the multi-delayed Lead-Lag transformations.




Figure 22: Illustration of multi-delayed lead-lag transformation. The dimension of lead-lag paths is d_{LLT} . The delayed paths are padded with zeros to ensure a fixed length for each dimension - Figure 2 in [JLNSY17]. Reproduced under license from Springer Nature (License Number 5425270874131).

The spatial features extracted by PSF_2 with truncation level n_{SP} of the paths corresponding to the pairs of joints in each frame (i.e. $m = 2$) are labelled **S-P-PSF**. The spatial features extracted by the PSF_3 with truncation level n_{ST} of the paths corresponding to the triples of joints in each frame (i.e. $m = 3$) are labelled **S-T-PSF**. The truncation levels $n_{SP} := 2$ and $n_{ST} := 3$ was found to be optimal in Section 6 of [JLNSY17].

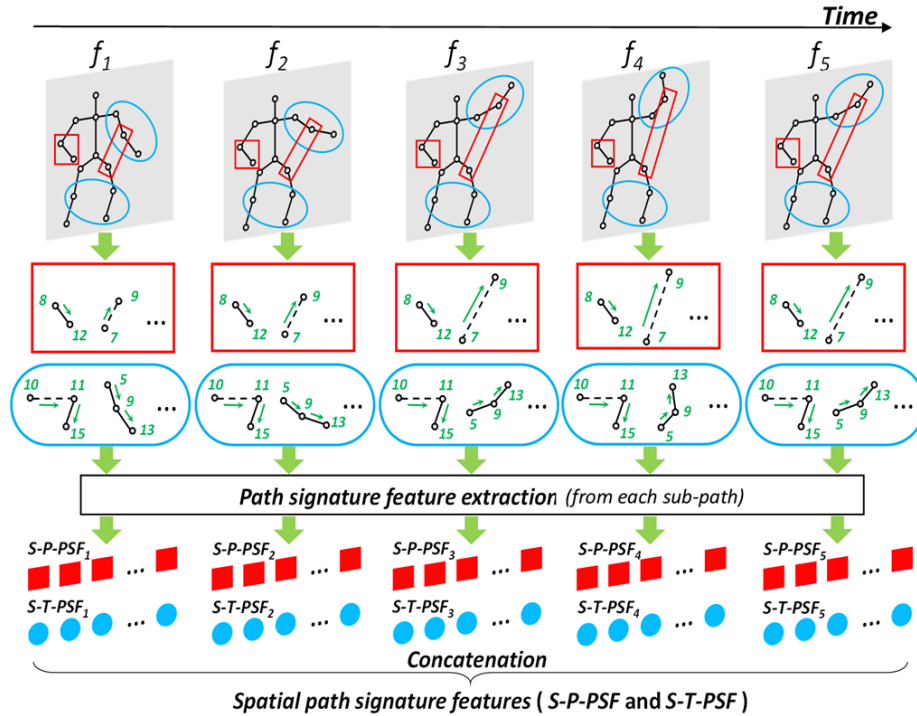


Figure 23: The illustration of spatial feature (**S-P-PSF** and **S-T-PSF**) extraction ($N=15$ in this figure). The red quadrangles denote the feature extraction of joint pairs, while the blue ellipses denote that of joint triples. All possible pairs and triples of joints are considered - Figure 3 in [JLNSY17]. Reproduced under license from Springer Nature (License Number 5425270874131).

The temporal features extracted by $\mathcal{D}_{h_{TJ}} - PSF_1$ with truncation level n_{TJ} of the paths corresponding to the evolution of the joint locations (i.e. $m = 1$) are labelled **T-J-PSF**. Each dimension of the features **S-P-PSF** and **S-T-PSF** characterises one particular spatial constraint of a pose. The temporal features extracted by $\mathcal{D}_{h_{TS}} - PSF_1$ with truncation level n_{TS} of multi-delay lead-lag transformations, with no-future-pause, of the one-dimensional paths corresponding to the evolution of each particular spatial constraint of a pose are labelled **T-S-PSF**. The choices of truncation levels $n_{TJ} := 5$ and $n_{TS} := 2$ and the choices of the hierarchical levels $h_{TJ} := 3$ and $h_{TS} := 3$ for the dyadic decomposition level were found to be optimal in Section 6 of [JLNSY17].

Illustrations of the extraction of the temporal features **T-J-PSF** and **T-S-PSF** are provided in Figures 24 and 25 respectively.




Figure 24: Illustration of temporal features extracted from the evolution of each corresponding joint (**T-J-PSF**) - Figure 4 in [JLNSY17]. Reproduced under license from Springer Nature (License Number 5425270874131).

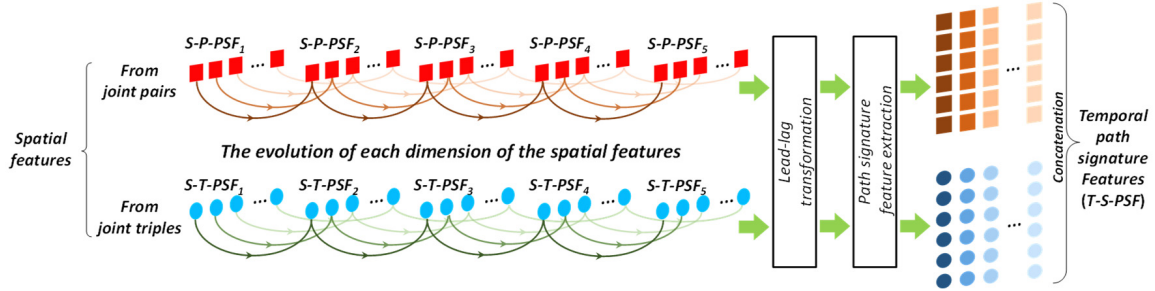


Figure 25: Illustration of temporal features extracted from the evolution of spatial context (**T-S-PSF**). Each dimension of the spatial features is treated equally and individually - Figure 5 in [JLNSY17]. Reproduced under license from Springer Nature (License Number 5425270874131).

Finally, the d -dimensional coordinates of the joint locations at each frame are recorded and labelled **S-J**. The features considered in [JLNSY17] are summarised below, starting with the **Spatial Structural Features** in each frame.

- **S-J**: The d -dimensional coordinates of each of the joints
- **S-P-PSF**: The PSF over each pair of joints up to signature level n_{SP}
- **S-T-PSF**: The PSF over each triple of joints up to signature level n_{ST}

The **Temporal Dynamical Features** in each frame are as follows.

- **T-J-PSF**: The temporal PSF over the evolution of each joint up to signature level n_{TJ}

- **T-S-PSF**: The evolution of each dimension of spatial PSF is treated as a path over which the temporal PSF up to signature level n_{TS} is extracted

Having extracted the features, we now describe the proposed models architecture. A single-hidden-layer neural network is chosen for the classifier, with the hidden layer consisting of 64 neurons. The input dimension is determined by the signature transform (i.e. by **S-P-PSF**, **S-T-PSF**, **T-J-PSF** and **T-S-SPF**) and the output is a probability distribution given by a softmax layer over all the class labels in the dataset. Training uses stochastic gradient descent, with early stopping implemented to allow a maximum of 300 epochs, on the cross-entropy cost function with a learning rate that decays exponentially in the number of epochs run.

A random proportion of the connections between the input and the single hidden layer are omitted using Dropconnect, see [FLWZZ13], for the purposes of regularisation. The high dimension of the data means that a high ratio of Dropconnect is required to prevent over-fitting, and the rate is set at 0.95.

12.4. Comparison with Specifically-Tailored Methods

The models performance on the four datasets MHAD, SBU, JHMDB and NTURGB+D is compared the performance of several models developed via other non-signature based techniques in [JLNSY17]. The results obtained for each of these four datasets are summarised below.

For the JHMDB dataset the off-the-shelf pose estimation called Alphapose (with Poseflow) [FLT17] was used to obtain 17 estimated joints from the RGB videos. The signature method proposed in [JLNSY17] is compared with the skeleton-based methods DT-FV [SW13], HLPF [BGJSZ13], Novel HLPF [FTZ13], and P-CNN [CLS15], all trained using the evolution of these 17 joints, in [JLNSY17]. As recorded in [JLNSY17], the accuracy 80.4% achieved by the Path Signature method [JLNSY17] is better than the accuracy achieved by any of the other methods (65.9% for DT-FV [SW13], 74.6% for P-CNN [CLS15], 76.0% for HLPF [BGJSZ13], and 79.6% for Novel HLPF [FTZ13]).

For the SBU dataset the two human bodies were regarded as a single united system with a total of 30 joints in 3D. The signature method proposed in [JLNSY17] is compared with the skeleton-based methods Yun et al. [BCHSY12], Ji et al. [CJY14], CHARM [CLLW15], HBRNN [DWW15], Deep HBRNN [LLSXXZZ16], Co-occurrence [LLSXXZZ16], STA-LSTM [LLSXZ17], ST-LSTM-Trust Gate [LSWX16, CLSWX17], SkeletonNet [ABBKS17], and GC-Attention-LSTM [DHKLW17], all trained using the evolution of the 30 joints, in [JLNSY17]. As recored in Table 6 in [JLNSY17], the accuracy of 96.8% for the Path Signature method [JLNSY17] is better than the accuracy achieved by any of the other methods (80.3% for Yun et al. [BCHSY12], 86.8% for Ji et al. [CJY14], 83.9% for CHARM [CLLW15], 80.4% for HBRNN [DWW15], 86.0% for Deep HBRNN [LLSXXZZ16], 90.4% for Co-occurrence [LLSXXZZ16], 91.5% for STA-LSTM [LLSXZ17], 93.3% for ST-LSTM-Trust Gate [LSWX16, CLSWX17], 93.5% for SkeletonNet [ABBKS17], and 94.1% for GC-Attention-LSTM [DHKLW17]).

The joint locations are recorded precisely enough in the Berkely MHAD dataset to enable methods to achieve perfect accuracy. The signature method proposed in [JLNSY17] is compared with the skeleton-based methods Vantigodi et al. [BV13], Ofli et al. [WWY13], Vantigodi et al. [RV14], Kapsouras et al. [KN14], HBRNN [DWW15], and ST-LSTM-Trust Gate [LSWX16, CLSWX17] on the Berkeley MHAD dataset in [JLNSY17]. As recorded in Table 7 in [JLNSY17], the perfect accuracy of 100% achieved by the Path Signature method [JLNSY17] matches the perfect accuracy of 100% achieved by both the HBRNN [DWW15] and SL-LSTM-Trust Gate [LSWX16, CLSWX17] methods, and is better than the accuracy achieved by the remaining methods (96.1% for Vantigodi et al. [BV13], 95.4% for Ofli et al. [WWY13], 97.6% for Vantigodi et al. [RV14], and 98.2% for Kapsouras et al. [KN14]).

For the NTURGB+D dataset, experiments were conducted on recognising both the cross-subject (i.e. actions performed by different people) and the cross-view (i.e. actions recorded from different view points) tasks. This involved an initial classification step to sort the data into 1-body and 2-body actions, before extracting the joints to be used. For both experiments the signature method proposed in [JLNSY17] is compared with the skeleton-based methods Dynamic Skeletons [HLZZ15], HBRNN [DWW15], Part-aware LSTM [LNSW16], ST-LSTM-Trust Gate [LSWX16, CLSWX17], STA-LSTM [LLSXZ17], SkeletonNet [ABBKS17], Joint Distance Maps [HLLW17], and GC-Attention-LSTM [DHKLW17] in [JLNSY17]. As recored in Table 9 in [JLNSY17], for the cross-subject task recognition, the accuracy of 78.3% achieved by the Path Signature method [JLNSY17] is better than the accuracy achieved by any of the other methods (60.2% for Dynamic

Skeletons [HLZZ15], 59.1% for HBRNN [DWW15], 62.9% for Part-aware LSTM [LNSW16], 69.2% for ST-LSTM-Trust Gate [LSWX16, CLSWX17], 73.4% for STA-LSTM [LLSXZ17], 75.9% for SkeletonNet [ABBKS17], 76.2% for Joint Distance Maps [HLLW17], and 74.4% for GC-Attention-LSTM [DHKLW17]). Additionally, as recored in Table 9 in [JLNSY17], for the cross-view task recognition, the accuracy of **86.3%** achieved by the Path Signature method [JLNSY17] is better than the accuracy achieved by any of the other methods (65.2% for Dynamic Skeletons [HLZZ15], 64.0% for HBRNN [DWW15], 70.3% for Part-aware LSTM [LNSW16], 77.7% for ST-LSTM-Trust Gate [LSWX16, CLSWX17], 81.2% for STA-LSTM [LLSXZ17], 81.2% for SkeletonNet [ABBKS17], 82.3% for Joint Distance Maps [HLLW17], and 82.8% for GC-Attention-LSTM [DHKLW17]).

The experiments conducted in [JLNSY17] illustrate that, within the class of recognising actions via the evolution of connected segments between landmark locations, the Path Signature method proposed in [JLNSY17], a relatively simple linear shallow fully-connected neural network, achieves comparable results to other more complex models.

12.5. Demo Notebook

The Landmark Human Action Recognition python notebook, created by Peter Foster and Kevin Schlegel, is based on [JLNSY17]. It provides an introduction to the methodology of [JLNSY17], presenting python code that uses the JHMDB dataset to generate a feature set as described in Section 12.3, before training a simple classifier action classifier using this feature set (again following the process discussed in Section 12.3).

The learning uses the PyTorch python package. The correct version of PyTorch to install depends on the particular hardware and operating system used. Consequently it is highly recommended to install PyTorch manually following the official instructions. After installing PyTorch, following the notebook is straightforward and once complete, the learnt classifier achieves an accuracy of 67.54%. This is lower than the 80.4% reported in [JLNSY17] (see Table 5 in [JLNSY17]) since the notebook chose path transformations for the purpose achieving a high speed of computations (with the feature set generated within seconds and the network training completing within a few minutes) rather than for maximal accuracy. By changing the path transformations and combining feature vectors of different sets of transformations the accuracy can be increased.

13. Distribution Regression via the Expected Signature

Many real-world tasks fall within the *Distribution Regression* framework, where the aim is to learn the functional relationship between multiple time-series and a single output. Examples range from determining the temperature of a gas from the trajectories of particles [Hil86, Rei99, Sch89], to using high-resolution climatic data to predict overall end-of-year crop yields [APP10, DR14, ELLLY17], and estimating mean-reversion parameters from observed financial market dynamics [LP11, GJR18, BGW00].

In Section 6, an approach to distribution regression via the expected signature was outlined. This approach is implemented in [BDLLS20], and we will provide a more detailed summary of their methodology within this section. Following the notation used in [BDLLS20], we let $\mathcal{C}([a, b]; V)$ denote the set of paths resulting from the removal of all the tree-like (see [HL10]) paths from $\mathcal{V}_1([a, b]; V)$. This has no practical impact on the following strategies [BDLLS20]. Consider M input-output pairs $\left\{ \left\{ x^{i,j} \right\}_{j=1}^{N_i}, y^i \right\}_{i=1}^M$ where each pair i is given by a target $y^i \in \mathbb{R}$ and a collection of N_i d -dimensional time-series $x^{i,j} = \left\{ \left(t_1, x_1^{i,j} \right), \dots, \left(t_{l_{i,j}}, x_{l_{i,j}}^{i,j} \right) \right\}$ of lengths $l_{i,j} \in \mathbb{N}$, time stamps $t_1 \leq \dots \leq t_{l_{i,j}}$ and values in a d -dimensional Banach space V . After making modelling assumptions to decide how the discrete streams should be converted to continuous paths, we have M groups of input-output pairs of the form

$$\left(\left\{ x^{1,j} : [a, b] \rightarrow V \right\}_{j=1}^{N_1}, y^1 \in \mathbb{R} \right), \dots, \left(\left\{ x^{M,j} : [a, b] \rightarrow V \right\}_{j=1}^{N_M}, y^M \in \mathbb{R} \right). \quad (13.1)$$

The collection of trajectories in group i is summarised by considering the measure $\delta^i := \frac{1}{N_i} \sum_{j=1}^{N_i} \delta_{x^{i,j}} \in \mathcal{PC}([a, b]; V)$ where $\delta_{x^{i,j}}$ is the Dirac measure centred at $x^{i,j}$. The input-output pairs in (13.1) can be repre-

sented as

$$(\delta^1 \in \mathcal{P}(\mathcal{K}), y^1 \in \mathbb{R}), \dots, (\delta^M \in \mathcal{P}(\mathcal{K}), y^M \in \mathbb{R}) \quad (13.2)$$

and we aim to learn a function $F : \mathcal{PC}([a, b]; V) \rightarrow \mathbb{R}$ from the pairs in (13.2).

The expected signature can characterise probability measures on paths (cf. Theorem 6.1) and in [BDLLS20] the authors propose two methods for distribution learning using the expected signature. The first is the approach outlined at the end of Section 6. To provide the details of this approach, we first recall that the pathwise expected signature $\Phi_{\text{Path}} : \mathcal{PC}([a, b]; V) \rightarrow \mathcal{C}([a, b]; T(V))$ is defined by

$$\Phi_{\text{Path}}(\mu)_t := \mathbb{E}_\mu [S_{a,t}(X)]. \quad (13.3)$$

A consequence of the universality of the pathwise expected signature (cf. Theorem 6.1) is that for any compact subset $\mathcal{K} \subset \mathcal{C}([a, b]; V)$, if F denotes our target function we want to learn, then

$$F|_{\mathcal{P}(\mathcal{K})} \in \text{Span} \left\{ S_{a,b}(\Phi_{\text{Path}}(\mu))^{\mathbf{J}} : \mu \in \mathcal{P}(\mathcal{K}), \mathbf{J} = (j_1, \dots, j_k) \in \{1, \dots, k\}^d, k \in \mathbb{N}_0 \right\} =: \mathcal{A}(\mathcal{K}). \quad (13.4)$$

The SES method of [BDLLS20] uses a truncation of $\mathcal{A}(Z)$, where $Z := \{\delta^i : i \in \{1, \dots, M\}\}$, as the feature functions before applying linear regression to find the optimal combination with respect to the mean squared error (MSE). The truncation is in terms of the length of the multi-indices allowed, i.e. in (13.4) we restrict to $k \leq K_0$ for some $K_0 \in \mathbb{N}$ so that only the coordinate iterated integrals of order up to K_0 are used.

The second method introduced in [BDLLS20] combines the expected signature with a Gaussian kernel. Recall that the expected signature $\mathbb{S} : \mathcal{C}([a, b]; V) \rightarrow T(V)$ is defined by

$$\mathbb{S}(\mu) := \mathbb{E}_\mu [S_{a,b}(X)] = \bigoplus_{n=0}^{\infty} \mathbb{E}_\mu [S_{a,b}^n(X)]. \quad (13.5)$$

The KES model proposed in [BDLLS20] considers the kernel $k : \mathcal{PC}([a, b]; V) \times \mathcal{PC}([a, b]; V) \rightarrow \mathbb{R}$ defined by

$$k(\mu, \nu) := \exp \left(-\sigma^2 \|\mathbb{S}(\mu) - \mathbb{S}(\nu)\|_{T(V)}^2 \right) \quad (13.6)$$

for $\sigma > 0$. It is established that the kernel k is universal in the sense that if $\mathcal{K} \subset \mathcal{C}([a, b]; V)$ is compact, then the associated *Reproducing Kernel Hilbert Space* (RKHS) is dense in $C^0(\mathcal{P}(\mathcal{K}); \mathbb{R})$ (see Theorem 3.3 in [BDLLS20]).

To evaluate $k(\delta^i, \delta^j)$ for $i, j \in \{1, \dots, M\}$ we first compute

$$\begin{aligned} \|\mathbb{S}(\delta^i) - \mathbb{S}(\delta^j)\|_{T(V)}^2 &= \left\| \frac{1}{N_i} \sum_{k=1}^{N_i} \mathbb{E}_{\delta_{x^{i,k}}} [S_{a,b}(X)] - \frac{1}{N_j} \sum_{l=1}^{N_j} \mathbb{E}_{\delta_{x^{j,l}}} [S_{a,b}(X)] \right\|_{T(V)}^2 \\ &= \left\| \frac{1}{N_i} \sum_{k=1}^{N_i} S_{a,b}(x^{i,k}) - \frac{1}{N_j} \sum_{l=1}^{N_j} S_{a,b}(x^{j,l}) \right\|_{T(V)}^2 \\ &= E_{ii} + E_{jj} - 2E_{ij} \end{aligned}$$

where

$$E_{ij} := \frac{1}{N_i N_j} \sum_{k=1}^{N_i} \sum_{l=1}^{N_j} \langle S_{a,b}(x^{i,k}), S_{a,b}(x^{j,l}) \rangle_{T(V)} = \frac{1}{N_i N_j} \sum_{k=1}^{N_i} \sum_{l=1}^{N_j} K_{x^{i,k}, x^{j,l}}(a, b). \quad (13.7)$$

In (13.7), $K_{\cdot, \cdot}(\cdot, \cdot)$ denotes the signature kernel discussed in Section 8. For the readers convenience, we recall that if $x \in \mathcal{V}_1([a, b]; V)$ and $y \in \mathcal{V}_1([c, d]; V)$ then $K_{x,y}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{T(V)}$ for $s \in [a, b]$ and $t \in [c, d]$ (cf. (8.1)). The relatively simple numerical scheme developed in [CFLSY20] (see Section 8) allows the inner products in (13.7) to be computed by a simple call to any numerical PDE solver of choice [BDLLS20].

The KES model uses the maps $\mu \mapsto k(\mu, \delta^i)$ for $i \in \{1, \dots, M\}$ for feature functions before applying linear regression to optimise the linear combination with respect to the mean squared error (MSE).

In Section 5 of [BDLLS20], the performance of the KES and SES models are compared with the performance of DeepSet method [KPRSSZ17] and other existing kernel-based methods for a variety of tasks. The kernel-based techniques considered all correspond to the same general framework. Distributions μ are first mapped to a RKHS H_1 via the map $\mu \mapsto \int_X k_1(\cdot, x)\mu(dx)$ where $k_1 : X \times X \rightarrow H_1$ is the reproducing kernel for H_1 . A second kernel k_2 is then used for the regression step to approximate a function $F : H_1 \rightarrow \mathbb{R}$ via the minimisation of a loss function. A more detailed summary may be found in Section 4 of [BDLLS20], with the full details appearing in [GSSS07], [DFMS12] [Fla15], [GPSS16], [FLSS17]. The notation DR- k_1 refers to the model produced by choosing k_2 to be a Gaussian kernel.

The particular kernel-based models considered are DR-RBF, DR-Matern32 and DR-GA, where GA refers to the Global Alignment kernel for time-series from [BCMv07], and the definition of all three may be found in Appendix B of [BDLLS20]. Kernel Ridge Regression is used to train KES and DR- k_1 for all choices of k_1 whilst Lasso Regression is used for SES. All models are run 5 times on each task with the mean and standard deviation of the predictive MSE recorded. The hyperparameters of KES and DR- k_1 were selected by cross-validation on the training set of each run. Full details of the implementation can be found in appendix B of [BDLLS20].

The first experiment considered determining the temperature T of an ideal gas. The simulation modelled 50 different gases, each consisting of 20 particles, by randomly initialising all velocities and letting the particles evolve at constant speed. The task is to learn T (sampled uniformly at random from $[1, 1000]$) from the set of trajectories traced by the particles in the gas. The complexity of the large-scale dynamics depends on both T and the radius of the particles. The results of two experiments, one where particles have a small radius (few collisions) and another where they have a bigger radius (many collisions), are recorded.

Model	Predictive MSE - Few Collisions	Predictive MSE - Many Collisions
DeepSets	8.69 ± 3.74	5.61 ± 0.91
DR-RBF	3.08 ± 0.39	4.36 ± 0.64
DR-Matern32	3.54 ± 0.48	4.12 ± 0.39
DR-GA	2.85 ± 0.43	3.69 ± 0.36
KES	1.31 ± 0.34	0.08 ± 0.02
SES	(1.26 ± 0.23)	0.09 ± 0.03

Table 19: Inferring the temperature of an ideal gas. Predictive MSE have been scaled by a factor of 100; if X is a predictive MSE value appearing in this table, then the actual predictive MSE is X/100. Table 1 in [BDLLS20]

The performances are comparable in the simpler setting, with KES and SES being slightly better. In the more collisions setting the KES and SES models significantly out-perform all other methods.

The *fractional Ornstein-Uhlenbeck* (fOU) process $\sigma_t = \exp(P_t)$, with $dP_t = -a(P_t - m)dt + vdB_t^H$ for $a, m, v \geq 0$ and B_t^H denoting a fractional Brownian motion of Hurst parameter $H \in (0, 1)$, is used as a model for volatility [ASS20]. Due to findings in [GJR18], H is chosen to be 0.2. The authors consider the task of estimating the mean-reversion parameter a from simulated sample paths of σ_t in [BDLLS20]. For this purpose, 50 mean-reversion values $\{a_i\}_{i=1}^{50} \subset [10^{-6}, 1]$ are chosen uniformly, before each a_i is regressed on a collection of $N \in \{20, 50, 100\}$ (time-augmented) trajectories $\left\{\sigma_t^{i,j}\right\}_{j=1}^N$ of length 200.

Model	Predictive MSE		
	N=20	N=50	N=100
Deepsets	74.43 \pm 47.57	74.07 \pm 49.15	74.03 \pm 47.12
DR-RBF	52.25 \pm 11.20	58.17 \pm 19.05	44.30 \pm 7.12
DR-Matern32	48.62 \pm 10.30	54.91 \pm 12.02	32.99 \pm 5.08
DR-GA	3.17 \pm 1.59	2.45 \pm 2.73	0.70 \pm 0.42
KES	1.41 \pm 0.40	0.30 \pm 0.07	0.16 \pm 0.03
SES	1.49 \pm 0.39	0.33 \pm 0.12	0.21 \pm 0.05

Table 20: Estimating mean-reversion parameters. Predictive MSE have been scaled by a factor of 1000; if X is a predictive MSE value appearing in this table, then the actual predictive MSE is X/1000. Table 2 in [BDLLS20].

The models KES and SES out-perform the others and progressively improve with the number of samples N .

The final task selected in [BDLLS20] for KES and SES to be evaluated on is a crop yield prediction task. The goal is predicting the yield of wheat crops over a region from the longitudinal measurements of climatic variables recorded across different locations of the region. The Eurostat dataset containing the total annual regional yield of wheat crops in mainland France - divided in 22 administrative regions - from 2015 to 2017 is used. Climatic measurements (temperature, soil humidity and precipitation), recorded at a frequency of every 6 hours, were extracted for each region. The total number of recordings varies across regions, and half of each regions recordings were randomly discarded as a further subsampling step. In addition to the MSE, the *mean absolute percentage error* (MAPE) was recorded. The baseline method predicts the average yield calculated from the training set.

Model	MSE	MAPE
Baseline	2.38 \pm 0.60	23.31 \pm 4.42
DeepSets	2.67 \pm 1.02	22.88 \pm 4.99
DR-RBF	0.82 \pm 0.22	13.18 \pm 2.52
DR-Matern32	0.82 \pm 0.23	13.18 \pm 2.53
DR-GA	0.72 \pm 0.19	12.55 \pm 1.74
KES	0.65 \pm 0.18	12.34 \pm 2.32
SES	0.62 \pm 0.10	10.98 \pm 1.12

Table 21: Predicting wheat crop yield - Table 3 in [BDLLS20]

The models KES and SES achieve larger improvements over the baseline than the other methods, with the SES model providing the best performance.

14. Anomaly Detection

Anomaly detection is the task of determining whether a given observation is unusual compared to a corpus of observations deemed to be usual. It arises in numerous fields such as medicine [BCCHVV13], financial fraud [CLNTZ16] and cybersecurity [JS00]. A natural approach is to use a distance metric, and view an event as an anomaly if it is at least some distance from the set of observations. Approaches to unsupervised anomaly detection for multivariate data include density-based approaches [BKNS00], clustering [DHX03], random forests [LTZ12], support vector machines [AAG13], neural networks [CC19] and nearest neighbour-based approaches [FHK04, HW07, KS20]. The article [CCPT14] provides a survey of a wide-range of anomaly detection techniques.

It is important to fix a sensible notion of what determines an outlier. It may initially seem tempting to determine an outlier by its difference to *most* other instances; i.e. by comparison to some averaged quantity determined by the entire collection. However this viewpoint runs both the risk of incorrectly deeming a data point to be an outlier, and of incorrectly failing to deem a data point as an outlier.

For an example of incorrectly deeming a data point to be an outlier, consider the world population as a dataset and suppose we are interested in the time taken by a person to run 100m. If compared to the average

time, it is likely that an elite sprinter would be considered an outlier. But of course they are not an outlier since there are *other* elite sprinters exhibiting similar performance amongst the population.

For an example of incorrectly deeming a data point to not be an outlier, assume we have a dataset $\Omega \subset \mathbb{R}^2$. Let $S^1 := \{x \in \mathbb{R}^2 : |x| = 1\} \subset \mathbb{R}^2$ denote the unit-circle in \mathbb{R}^2 with respect to the Euclidean distance on \mathbb{R}^2 . Let $\varepsilon \in (0, 1)$ and suppose that $\Omega = \tilde{\Omega} \cup \{x_*\}$ where $\tilde{\Omega} \subset S^1$ is a uniformly distributed sample of points from S^1 , and $x_* \in \mathbb{R}^2 \setminus S^1$ with $\text{dist}(x_*, S^1) = \varepsilon$. Assume that the cardinality of $\tilde{\Omega}$ is sufficiently large and $\varepsilon > 0$ is sufficiently small so that the average position of the points in Ω is very close to the origin in \mathbb{R}^2 . Then in terms of absolute magnitude (i.e. size of norm) the point x_* will not be deemed an outlier. But being the *only* point in Ω to *not* lie within the unit-circle S^1 makes the point x_* different from *all* other points in Ω , and therefore it should be an outlier.

With this in mind, a more useful point-of-view is to determine outliers by *not* being similar to *any* other instance within the dataset. The *k-nearest neighbour* (k-NN) approach uses this idea by comparing an instance to its k nearest neighbours within the dataset. It is based on the assumption that usual data points have close neighbours in the training set, while unusual points are located far from their neighbours [FHK04]. Loosely, a point is declared an outlier if it is located far from its k-nearest neighbours. The choice of k can be treated as a hyperparameter to be optimised.

The k-NN approach requires one to make a choice of distance metric to use for the determination of a points k-nearest neighbours. The *Mahalanobis distance* [Mah36], providing a notion of distance between a point and a distribution, is a popular choice considered in the works [CSSQZ20, KS20], for example. In particular, the recent work [ACFL20] considers a combination of the Mahalanobis distance with k-NN techniques for the purpose of anomaly detection. The anomaly detection is implemented on a collection of paths; namely, it determines whether an entire path fails to be similar to any of the other paths within the collection. This is achieved in [ACFL20] via use of the Mahalanobis distance and k-NN techniques on a vector feature set determined by taking truncated signatures of the paths.

The remainder of this section is dedicated to covering the approach proposed in [ACFL20]. We closely follow the presentation of [ACFL20], beginning with the introduction of the Mahalanobis distance. In order to match with [ACFL20], we adopt their terminology of *variance* for the Mahalanobis distance (which is arguably more natural).

Given a probability measure μ on a vector space V , denoted $\mu \in \mathcal{P}(V)$, the covariance quadratic form $\text{Cov}_\mu : V^* \times V^* \rightarrow \mathbb{R}$ is given by $\text{Cov}_\mu(\phi, \psi) := \mathbb{E}_\mu[\phi(\cdot)\psi(\cdot)]$. It induces a dual norm defined for $v \in V$ by

$$\|v\|_\mu := \sup \{\phi(v) : \text{Cov}_\mu(\phi, \phi) \leq 1\}. \quad (14.1)$$

It is finite within the linear span of the support of μ and infinite outside. This norm, computed for the measure μ re-centred to have mean zero, is the *variance norm* associated to μ . The *conformance* of $x \in V$ with respect to $\mu \in \mathcal{P}(V)$ is defined to be the distance

$$\text{dist}(x; \mu) := \inf \left\{ \|x - y\|_\mu : y \in \text{support}(\mu) \right\}. \quad (14.2)$$

We now motivate the usefulness of the variance norm for anomaly detection via consideration of the Gaussian setting as done in [ACFL20].

Let $Z = (z_1, \dots, z_d)$ be the standard normal distribution on \mathbb{R}^d , where the z_i are independent normal variables of mean zero and variance 1 on some probability space $(\mathcal{N}_i, \mathbb{P}_i)$. Let γ denote the probability measure induced on \mathbb{R}^d as the product of the push-forward measures $(z_i)_\# \mathbb{P}_i$ on \mathbb{R} . Of course this means nothing more than that γ is the standard Gaussian measure on \mathbb{R}^d defined by

$$\gamma(A) := (2\pi)^{-\frac{d}{2}} \int_A e^{-\frac{|x|^2}{2}} dx \quad (14.3)$$

for Borel-measurable sets $A \subset \mathbb{R}^d$.

The Riesz Representation Theorem tells us that for any $\phi \in (\mathbb{R}^d)^*$ there is a unique $a(\phi) \in \mathbb{R}^d$ for which $\phi(\cdot) \equiv \langle \cdot, a(\phi) \rangle$ where $\langle \cdot, \cdot \rangle$ denotes the standard Euclidean inner product on \mathbb{R}^d . A computation tells us that the covariance is given by the standard Euclidean inner product in the sense that $\text{Cov}_\gamma(\phi, \psi) = \langle a(\phi), a(\psi) \rangle$

whenever $\phi, \psi \in (\mathbb{R}^d)^*$. Thus the dual norm is the usual Euclidean norm

$$\|v\|_\gamma := \sup \{ \phi(v) : \text{Cov}_\gamma(\phi, \phi) \leq 1 \} = \sup \{ \langle v, a \rangle : \langle a, a \rangle \leq 1 \} = |v|.$$

If we consider $|Z|^2$ then its expected value is d , and so the norm of the path is of order \sqrt{d} . For high dimensions d the norm is huge and in the infinite dimensional case it is infinite. For Brownian motion B_t on $[0, T]$ it is the L^2 norm of the gradient, and clearly no Brownian path has a derivative, let alone a square integrable one. Whilst this notion of variance is intrinsic, it appears to demand a notion of closeness that is completely unrealistic.

Yet this notion of variance provides an excellent measure of exceptionality. That this is true is a consequence of the *Tsirelson-Sudakov-Borell (TSB) Isoperimetric Inequality* [AT07]. The TSB Isoperimetric Inequality itself asserts that half-spaces have minimal Gaussian boundary measure among all sets of a fixed Gaussian measure.

Theorem 14.1 (TSB Isoperimetric Inequality). *Let $A \subset \mathbb{R}^d$ be γ -measurable and for $\varepsilon > 0$ let A_ε denote the ε -fattening of A , i.e. $A_\varepsilon := \{x \in \mathbb{R}^d : \text{dist}(x, A) \leq \varepsilon\}$. Define $\varphi, \Phi : \mathbb{R} \rightarrow \mathbb{R}$ by*

$$\varphi(x) := \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \quad \text{and} \quad \Phi(x) := \int_{-\infty}^x \varphi(s) ds. \quad (14.4)$$

Then

$$\liminf_{\varepsilon \downarrow 0} \left(\frac{\gamma(A_\varepsilon) - \gamma(A)}{\varepsilon} \right) \geq \varphi(\Phi^{-1}(\gamma(A))) \quad (14.5)$$

with equality in (14.5) if and only if A is a half-space. Consequently, whenever H is a half-space with $\gamma(H) = \gamma(A)$, for every $r \geq 0$ we have

$$\gamma(A_r) \geq \gamma(H_r). \quad (14.6)$$

To illustrate the variances measure of exceptionality, suppose that a subset A satisfies that $\gamma(A) = \frac{1}{2}$. Let H denote a half-space with $\gamma(H) = \gamma(A) = \frac{1}{2}$. The rotational symmetry of γ means we may assume that $H = \mathbb{R}^{d-1} \times (-\infty, 0]$ so that $\gamma(A) = \gamma(H) = \Phi(0)$. Then $H_r = \mathbb{R}^{d-1} \times (-\infty, r)$ and $\gamma(H_r) = \Phi(r)$. Consequently, the TSB Isoperimetric Inequality tells us that (cf. (14.6))

$$\gamma(A_r) \geq \gamma(H_r) = \Phi(r). \quad (14.7)$$

Therefore if we choose a new instance of Z and label it Z_{new} , then

$$\mathbb{P}[\{\text{dist}(Z_{\text{new}}, A) \geq r\}] = 1 - \gamma(A_r) \stackrel{(14.7)}{\leq} 1 - \Phi(r) = \frac{1}{\sqrt{2\pi}} \int_r^\infty e^{-\frac{w^2}{2}} dw. \quad (14.8)$$

The integral in (14.8) becomes arbitrarily small for even moderately sized $r > 0$, meaning that it is very unlikely for the new instance Z_{new} to be far away from the entirety of A . Returning to the Brownian motion example, this reflects that whilst a Brownian path B_t may be irregular, if you select a corpus of Brownian paths with probability $1/2$, then B_t will differ from one of those paths by a differentiable path of small norm.

These properties of the variance are desirable from the perspective of measuring exceptionality; it is selective and discriminatory, but it must be used to compare with the corpus and not directly. A new element of the corpus will be far away from most member, but with very high probability there will be some members with which it is associated well.

Having motivated the variances suitability for anomaly detection, we return to considering a collection of streams. Empirical measures are natural to use to summarise a collection of streams (cf. Section 13). However such measures will never be Gaussian. Consequently, taking half the ensemble will only adequately capture the other half if the sample size is large enough. In this context, large enough means that the balls of a certain radius centred at each point capture a good proportion of the measure. Without reaching this threshold, essentially every sample will be non-conformant.

The authors use k-NN techniques to develop an empirical route around this issue in [ACFL20]. If the corpus is randomly split into two halves, then we may determine R such that, with probability $1/2$, given any point x in the second half of the corpus, its nearest neighbour y within the first half satisfies that $\|x - y\|_\mu \leq R$.

This R is referred to as the *threshold of conformance* of the corpus. From this scale on, if a new observation to the corpus has conformance $r + R$, then r should provide an effective measure of being an anomaly whilst R provides a measure of the extent to which the feature set is overwhelming the sample size. This empirical approach can be summarised as follows.

- i. Apply the truncated signature transform $\Pi_N(S)$ (cf. Section 7 for how to choose $N \in \mathbb{N}$) to map the given finite corpus of data $\mathcal{D} \subset V$ to a corpus $\mathcal{C} \subset T^{(N)}(V)$. As mentioned above, \mathcal{C} has its own threshold of conformance, denoted $R_{\mathcal{C}}$.
- ii. Let $\text{Sig}^N(\mathcal{C})$ denote the normalised sum of dirac delta measures that is empirically determined by $\text{Sig}^N(\mathcal{C})$, i.e. place a dirac mass at each point in \mathcal{C} determined by the signature, and take the sum normalised to have total mass 1.
- iii. Randomly split $\mathcal{C} = \mathcal{C}_1 \sqcup \mathcal{C}_2$ and let μ_1 and μ_2 denote the restriction of $\text{Sig}^N(\mathcal{C})$ to \mathcal{C}_1 and \mathcal{C}_2 respectively.
- iv. For a random point $x \in V$ with law μ_1 we consider the random variable $\text{dist}(x; \mu_1)$. For $\varepsilon > 0$, the region where $\mu_2(\text{dist}(x; \mu_1) \geq R_{\mathcal{C}}) = \varepsilon$ yields a natural quantified choice of anomalous behaviour.

Computing the variance norm $\|w\|_{\text{Sig}^N(\mathcal{C})}$ requires knowledge of the expected signature up to level $2N$, see Propositions 3.1 and 3.2 in [ACFL20]. Hence this strategy requires us to compute the signature transform truncated to order $2N$. Intuitively, the order N is a measure of the resolution at which the streams are viewed. For small N , only general features of the streams are considered, with more details of the streams considered as N increases.

In [ACFL20] it is established that the variance norm measures anomalous behaviour in a sensible way. Firstly, the extent to which a stream outside the corpus is considered an outlier increases as the order N increases (cf. Proposition C.1 in [ACFL20]). Additionally, the variance of any stream not in the corpus is infinite provided we take a sufficiently high order resolution N (cf. Proposition C.2 in [ACFL20]).

Three major benefits of the method proposed in [ACFL20] are:

- The variance norm is intrinsic and independent of any choice of basis,
- The conformance score does not depend on any external choice of parameters (metric etc), and
- There are no distribution assumptions made on the corpus of data.

The approach is validated in [ACFL20] through several experiments for the task of unsupervised anomaly detection. For these experiments, we have a data set $\mathcal{I} \subset \mathcal{S}(\mathbb{R}^d)$ partitioned into data deemed to be normal $\mathcal{I}_{\text{Normal}}$ and those data deemed to be anomalous $\mathcal{I}_{\text{Anomaly}}$. Further partitioning provides the corpus $\mathcal{C} \subset \mathcal{I}$ which is used for training, with $\mathcal{Y} := \mathcal{I} \setminus \mathcal{C}$ providing the testing data.

The first experiment considers handwritten digits. The PenDigitis-orig data set [DG19] consisting of 10992 instances of handwritten digits captured from 44 subjects. Each instance is represented as a two-dimensional stream, based on sampling of the pen position. $\mathcal{I}_{\text{Normal}}$ is the set of digits labelled m , \mathcal{C} is the subset of $\mathcal{I}_{\text{Normal}}$ labelled as “training” by the annotators, and \mathcal{Y} is the set of instances labelled as “testing”. Finally, $\mathcal{I}_{\text{Anomaly}}$ is the subset of \mathcal{Y} not representing digit m . Signatures of order $N \in \{1, \dots, 5\}$ are considered without any augmentation, and the results are summarised in Table 22.

$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$
0.901 ± 0.004	0.965 ± 0.002	0.983 ± 0.001	0.987 ± 0.001	0.989 ± 0.000

Table 22: Handwritten digits data: performance quantified using ROC AUC in response to signature order N - Table 1 in [ACFL20]

Another experiment in [ACFL20] considers a sample of marine vessel traffic data based on the the automatic identification system (AIS) which reports a ship’s geographical position alongside other vessel information. They consider the stream of timestamped latitude/longitude position data associated with each vessel.

Vessel data with invalid identifiers or invalid length information were discarded. Streams are compressed to only contain points that are at least 10m away from the previous position, and vessels whose starting and finishing points are within 5km are discarded. To evaluate the effect of the stream length, the streams are disintegrated so that the length between initial and final points in each sub-stream remains constant with $D \in \{4\text{km}, 8\text{km}, 16\text{km}, 32\text{km}\}$. Only sub-streams whose maximum distance between successive points is less than 1km are subsequently retained. A sub-stream is deemed normal if it belongs to a vessel of length at least 100m, and deemed anomalous if it corresponds to a vessel of length no greater than 50m.

An order $N = 3$ truncated signature transform was then utilised to implement the anomaly detection method on the resulting corpus with a variety of augmentation transformations considered, full details of the exact implementation may be found in [ACFL20]. For comparison, as a baseline approach the authors summarise each sub-stream by estimating its component-wise mean and covariance, retaining the upper triangular part of the covariance matrix. This results in feature vectors of dimensionality $\frac{1}{2}(n^2 + 3n)$ which are provided as input to an isolation forest [LTZ08]. The isolation forest is trained using 100 trees and for each tree in the ensemble using 256 samples represented by a single random feature.

The following table displays results for the proposed signature approach in comparison to the baseline, for combinations of stream transformations and values of the sub-stream length D . In table 23, $\text{Sig}^3(\mathcal{C})$ denotes the signature conformance method proposed in [ACFL20] whilst Isolation denotes the isolation forest method used as a baseline.

Model	Lead-Lag	Time-Diff	Time-inv	D = 4km	D = 8km	D = 16km	D = 32km
Isolation	No	No	No	0.690	0.718	0.717	0.733
$\text{Sig}^3(\mathcal{C})$	No	No	No	0.723	0.706	0.705	0.740
Isolation	No	No	Yes	0.682	0.698	0.714	0.716
$\text{Sig}^3(\mathcal{C})$	No	No	Yes	0.776	0.789	0.785	0.805
Isolation	No	Yes	No	0.771	0.779	0.779	0.803
$\text{Sig}^3(\mathcal{C})$	No	Yes	No	0.810	0.813	0.818	0.848
Isolation	No	Yes	Yes	0.745	0.751	0.761	0.797
$\text{Sig}^3(\mathcal{C})$	No	Yes	Yes	0.839	0.860	0.863	0.879
Isolation	Yes	No	No	0.759	0.765	0.766	0.763
$\text{Sig}^3(\mathcal{C})$	Yes	No	No	0.811	0.835	0.824	0.837
Isolation	Yes	No	Yes	0.755	0.761	0.763	0.762
$\text{Sig}^3(\mathcal{C})$	Yes	No	Yes	0.812	0.835	0.833	0.855
Isolation	Yes	Yes	No	0.820	0.815	0.823	0.817
$\text{Sig}^3(\mathcal{C})$	Yes	Yes	No	0.845	0.861	0.862	0.877
Isolation	Yes	Yes	Yes	0.810	0.795	0.816	0.815
$\text{Sig}^3(\mathcal{C})$	Yes	Yes	Yes	0.848	0.863	0.870	0.891

Table 23: Marine vessel traffic data: performance quantified using ROC AUC - Table 2 in [ACFL20]

Signature conformance out-performs the baseline for 30 out of the 32 parameter combinations. The maximum ROC-AUC score (highlighted in bold in Table 23) is achieved by the signature conformance for $D = 32\text{km}$ using a combination of lead-lag, time differences and invisibility reset transformations. The score of 0.891 is a 6.8 percentage points gain compared to the best-performing baselines parameter combination.

References

- [AAG13] S. Abdennadher, M. Amer and M. Goldstein, *Enhancing one-class support vector machines for unsupervised anomaly detection*, In Proceedings of the ACM SIGKDD Workshop on Outlier Detection and Description, pages 8–15, 2013.
- [ABDGKKMMRRS15] Ahmed Abdulkadir, Meritxell Bach-Cuadra, Alessandro Daducci, Cristina Granziera, Stefan Klöppel, Gunnar Krueger, Philippe Maeder, Bénédicte Maréchal, Reto Meuli, Delphine

- Ribes, Alexis Roche and Daniel Schmitter, *An evaluation of volume-based morphometry for prediction of mild cognitive impairment and Alzheimer's disease*, *NeuroImage Clinical* **7**, 7-17, 2015.
- [AT07] R. J. Adler and J. E. Taylor, *Gaussian inequalities*, *Random Fields and Geometry*, pages 49–64, 2007.
- [AP17] Z. Aldeneh and E. M. Provost, *Using regional saliency for speech emotion recognition*, *ICASSP, IEEE*, pp. 2741-2745, 2017.
- [ABJMMNWW19] Emily Alsentzer, William Boag, Di Jindi, Matthew McDermott, John Murray, Tristan Naumann and Wei-Hung Weng, *Publicly available clinical bert embeddings*, In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72-78, 2019.
- [AHR20] Jaan Altosaar, Kexin Huang and Rajesh Ranganath, *Clinicalbert: Modeling clinical notes and predicting hospital readmission*, In *Proc. ACM Conference on Health, Inference and Learning (CHIL)*, 2020.
- [APP10] Daniel P Ames, Sudhanshu Sekhar Panda and Suranjan Panigrahi, *Application of vegetation indices for agricultural crop yield prediction using neural network techniques*, *Remote Sensing*, **2**(3):673–696, 2010.
- [ABBK17] S. An, M. Bennamoun, F. Boussaid, Q. Ke and F. Sohel, *SkeletonNet: mining deep part features for 3D action recognition*, *IEEE Signal Processing Letters*, vol. **24**, no. 6, pp. 731-735, 2017.
- [ABDEELNSSST16] E. André, C. Busso, L. Y. Devillers, J. Epps, F. Eyben, P. Laukka, S. S. Narayanan, K. R. Scherer, B. W. Schuller, J. Sundberg and K. P. Truong, *The Geneva Minimalistic Acoustic Parameter Set (GeMAPS) for Voice Research and Affective Computing*, *IEEE Transactions on Affective Computing*, **7**(2), pp. 190-202, 2016.
- [ABBCCDDFLS94] A. Andreoli, P. Berger, P. Buchheim, S. Channabasavanna, B. Coid, A. Dahl, R. F. Diekstra, B. Ferguson, A. W. Loranger, N. Sartorius et al., *The international personality disorder examination: The world health organization/alcohol, drug abuse, and mental health administration international pilot study of personality disorders*, *Archives of General Psychiatry*, 1994.
- [ABILNPS17] Cecilie Anker, Ioana Balas, Christian Igel, Martin Lillholm, Mads Nielsen, Akshay Pai and Lauge Sørensen. *Differential diagnosis of mild cognitive impairment and Alzheimer's disease using structural MRI cortical thickness, hippocampal shape, hippocampal texture, and volumetry*, *NeuroImage Clinical* **13**, 470-481, 2017.
- [ACV14] F. Arrate, R. Chellappa and R. Vemulapalli, *Human action recognition*, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp.588-595, 2014.
- [ABKLS19] Imanol Perez Arribas, Patric Bonnier, Patrick Kidger, Terry Lyons and Christopher Salvi, *Deep Signature Transforms*, *NeurIPS* 2019.
- [ACFL20] Imanol Perez Arribas, Thomas Cochrane, Peter Foster and Terry Lyons, *Anomaly detection on streamed data*, Preprint 2020. <https://arxiv.org/abs/2006.03487>
- [AGGLS18] Imanol P. Arribas, Guy M. Goodwin, John R. Geddes, Terry Lyons and Kate E. Saunders, *A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder*, *Translational psychiatry*, vol. **8**, no. 1, pp. 1-7, 2018.
- [ASS20] Imanol Perez Arribas, Cristopher Salvi and Lukasz Szpruch, *Sig-sdes model for quantitative finance*, *Proceedings of the First ACM International Conference on AI in Finance* 2020.
- [BV13] R. V. Babu and S. Vantigodi, *Real-time human action recognition from motion capture data*, In *National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics*, pp. 1-4, 2013.
- [BBDFKLLS18] A. Bagnall, A. Bostrom, H. A. Dau, M. Flynn, E. Keogh, J. Large, J. Lines and P. Southam, *The UEA multivariate time series classification archive*, *arXiv preprint*, 2018. <https://arxiv.org/abs/1811.00075>.

- [[BBKLL17](#)] A. Bagnall, A. Bostrom, E. Keogh, J. Large and J. Lines, *The Great Time Series Classification Bake Off: A Review and Experimental Evaluation of Recent Algorithmic Advances*, Data Mining and Knowledge Discovery, **31** (2017), pp.606-660, 2017.
- [[BMW98](#)] C. Bailer-Jones, D. MacKay, and P. Withers, *A recurrent neural network for modelling dynamical systems*, Network: Computation in Neural Systems, vol. **9**, pp.531-47, 1998
- [[BCKOV13](#)] R. Bajcsy, R. Chaudhry, G. Kurillo, F. Ofli and R. Vidal, *Berkeley MHAD: A comprehensive multimodal human action database*, In IEEE Workshop on Applications of Computer Vision (WACV), pp. 53-60, 2013.
- [[BGW00](#)] Ronald Balvers, Erik Gilliland and Yangru Wu, *Mean reversion across national stock markets and parametric contrarian investment strategies*, The Journal of Finance, **55**(2):745–772, 2000.
- [[BMZ17](#)] E. Barsoum, S. Mirsamadi and C. Zhang, *Automatic speech emotion recognition using recurrent neural networks with local attention* ICASSP 2017, IEEE, pp. 2227-2231, 2017.
- [[BCCHVV13](#)] I. Batal, G. Clermont, G. F. Cooper, M. Hauskrecht, M. Valko and S. Visweswaran, *Outlier detection for patient monitoring and alerting*, Journal of Biomedical Informatics, **46**(1):47–55, 2013.
- [[BCDFHS00](#)] D. W. Bednarski, N. Cristianini, N. Duffy, D. Haussler, T. S. Furey and M. Schummer, *Support Vector Machine Classification and Validation of Cancer Tissue Samples Using Microarray Expression Data*, Bioinformatics, **16**, pp.906-914, 2000.
- [[BCDR18](#)] J. Bettencourt, R. T. Q. Chen, D. K. Duvenaud and Y. Rubanova, *Neural Ordinary Differential Equations*, Advances in Neural Information Processing Systems **31**, pp.6571-6583, Curran Associates, Inc., 2018.
- [[BCHSY12](#)] T. L. Berg, D. Chattopadhyay, J. Honorio, D. Samaras and K. Yun, *Two-person interaction detection using body-pose features and multiple instance learning*, In Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 28-35, 2012.
- [[BCDOPSTV16](#)] A. C. Bilderbeck, G. M. Clifford, G. D. an Goodwin, M. Osipov, N. Palmius, K. E. A. Saunders, A. Tsanas and M. De Vos, *Daily longitudinal self-monitoring of mood variability in bipolar disorder and borderline personality disorder*, Journal of Affective Disorders, **205**:225-233, 2016.
- [[BCMV07](#)] O. Birkenes, M. Cuturi, T. Matsui and J.-P. Vert, *A Kernel for Time Series Based on Global Alignments*, In 2007 IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP07, volume **2**, pages II-413, IEEE, 2007
- [[BLNW20](#)] John Pougué Biyong, Terry Lyons, Alejo J Nevado-Holgado and Bo Wang, *Information Extraction from Swedish Medical Prescriptions with Sig-Transformer Encoder*, Proceedings of the 3rd Clinical Natural Language Processing Workshop, pages 41-54, 2020.
- [[BGJSZ13](#)] M. J. Black, J. Gall, H. Jhuang, C. Schmid and S. Zuffi, *Towards understanding action recognition*, In IEEE International Conference on Computer Vision (ICCV), pp. 3192-3199, 2013.
- [[BCFFKMSS13](#)] A. Blake, M. Cook, M. Finocchio, A. Fitzgibbon, A. Kipman, R. Moore, T. Sharp and J. Shotton, *Real-time human pose recognition in parts from single depth images*, Communications of the ACM, vol. **56**, no. 1, pp. 116-124, 2013.
- [[BDMN19](#)] Horatio Boedihardjo, Joscha Diehl, Marc Mezzarobba and Hao Ni, *The expected signature of Brownian motion stopped on the boundary of a circle has finite radius of convergence*, Bulletin of the London Mathematical Society, Volume **53**, Issue 1, 2021.
- [[BDLLS20](#)] Edwin V. Bonilla, Theodoros Damoulas, Maud Lemerrier, Terry Lyons and Cristopher Salvi, *Distribution Regression for Sequential Data*, Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume **130**, p.3754-3762, 2021.

- [BO19] Patric Bonnier and Harald Oberhauser, *Signature Cumulants, Ordered Partitions, and Independence of Stochastic Processes*,
- [BGLY14] Youness Boutaib, Lajos Gergely Gyurkó, Terry Lyons and Danyu Yang, *Dimension-free Euler estimates of rough differential equations*, *Revue Roumaine de Mathématiques Pures et Appliquées*, **59**, 2014.
- [BKNS00] M. M. Breunig, H. P. Kriegel, R. T. Ng, and J. Sander, *LOF: Identifying density-based local outliers*, In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, pages 93–104, 2000.
- [BSS20-I] T. G. Buchman, K. L. Sciarretta and S. Q. Simpson, *Sepsis among medicare beneficiaries: 1. The burdens of sepsis, 2012-2018*, *Crit Car, Med* 2020; **48**:276-288.
- [BSS20-II] T. G. Buchman, K. L. Sciarretta and S. Q. Simpson, *Sepsis among medicare beneficiaries: 2. The trajectories of sepsis, 2012-2018*, *Crit Car, Med* 2020; **48**:289-301.
- [BSS20-III] T. G. Buchman, K. L. Sciarretta and S. Q. Simpson, *Sepsis among medicare beneficiaries: 3. The methods, models, and forecasts of sepsis*, *Crit Car, Med* 2020; **48**:302-318.
- [BBCKKLLMN08] M. Bulut, C. Busso, J. N. Chang, S. Kim, C. C. Lee, S. Lee, E. Mower and S. S. Narayanan, *Iemocap: Interactive emotional dyadic motion capture database*, *Language resources and evaluation*, vol. **42**, no. 4, p. 335, 2008.
- [CFLSY20] Thomas Cass, James Foster, Terry Lyons, Christopher Salvi and Weixin Yang, *The Signature Kernel is the solution of a Goursat PDE*, *SIAM Journal on Mathematics of Data Science*, Vol. **3**, Iss. 3, p.873-899, 2021. <https://doi.org/10.1137/20M1366794>
- [CLX21] Thomas Cass, Terry Lyons and Xingcheng Xu, *General Signature Kernels*, arXiv preprint, 2021. <https://arxiv.org/abs/2107.00447>
- [CC19] R. Chalapathy and S. Chawla, *Deep learning for anomaly detection: A survey*, arXiv preprint, 2019. <https://arxiv.org/abs/1901.03407>
- [CDLT18] Ming-Wei Chang, Jacob Devlin, Kenton Lee and Kristina Toutanova, *Bert: Pre-training of deep bidirectional transformers for language understanding*, arXiv preprint, arXiv:1810.04805, 2018.
- [CT01] E. Change and S. Tong, *Support Vector Machine Active Learning for Image Retrieval*, In *Proceedings of the ninth ACM international conference on Multimedia*, pp.107-118, 2001.
- [CFGKKLPRSSTWZ06] M. Cheang, D. Feinstein, D. Gurka, Anand Kumar, Aseem Kumar, B. Light, J. E. Parrillo, D. Roberts, S. Sharma, R. Suppes, L. Taiberg, K. E. Wood and S. Zanotti, *Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock*, *Crit Care Med* 2006, **34**:1589-1596.
- [CDLW20] R. T. Q. Chen, D. Duvenaud, X. Li and T.-K. L. Wong, *Scalable Gradients and Variational Inference for Stochastic Differential Equations*, *AISTATS*, 2020.
- [CDR19] R. T. Chen, D. Duvenaud and Y. Rubanova, *Latent odes for irregularly-sampled time series*, *Advances in Neural Information Processing Systems* **32** (NeurIPS 2019).
- [CHLRSW18] X. Chen, W. Han, H. Li, H. Ruan, B. Schuller and Z. Wang, *Towards temporal modelling of categorical speech emotion recognition*, *Proc. Interspeech* 2018, pp. 931-936, 2018.
- [CSSQZ20] Zheng Chen, Yulin Si, Jili Sun, Peng Qian and Dahai Zhang, *A Data-Driven Fault Detection Framework Using Mahalanobis Distance Based Dynamic Time Warping*, *IEEE Access*, Volume **8**, June 2020.
- [CGGLLNPTU20] Hao Cheng, Jianfeng Gao, Yu Gu, Xiaodong Liu, Michael Lucas, Tristan Naumann, Hoi-fung Poon, Robert Tinn and Naoto Usuyama, *Domain-specific language model pretraining for biomedical natural language processing*, arXiv preprint, arXiv:2007.15779, 2020.

- [CJY14] H. Cheng, Y. Ji and G. Ye, *Interactive body part contrast minimizing for human interaction recognition*, In International Conference on Multimedia and Expo Workshops, pp. 1-6, 2014.
- [CLS15] G. Chéron, I. Laptev and C. Schmid, *P-cnn: Pose-based cnn features for action recognition*, In IEEE International Conference on Computer Vision (ICCV), pp. 3218-3226, 2015.
- [Che13] Ilya Chevyrev, *Unitary representations of geometric rough paths*,
- [CK16] Ilya Chevyrev and Andrey Kormilitzin, *A Primer on the Signature Method in Machine Learning*, arXiv preprint, 2016. <https://arxiv.org/abs/1603.03788>
- [CL13] Ilya Chevyrev and Terry Lyons, *Characteristic Functions of Measures On Geometric Rough Paths*, The Annals of Probability, Vol. **44**, No. 6, 4049-4082, 2016.
- [CO18] Ilya Chevyrev and Harald Oberhauser, *Signature moments to characterise laws of stochastic processes*, Preprint 2018. <https://arxiv.org/abs/1810.10971>
- [CCPT14] D. A. Clifton, L. Clifton, M. A. Pimentel and L. Tarassenko, *A review of novelty detection*, Signal Process., **99**, pp.215-249, 2014.
- [CGGMO18] M. Ciccone, M. Gallieri, F. Gomez, J. Masci and C. Osendorfer, *NAIS-Net: Stable Deep Networks from Non-Autonomous Differential Equations*, Advances in Neural Information Processing Systems **31**, pp.3025-3035, Curran Associates, Inc., 2018.
- [CLSWX17] A. K. Chichung, J. Liu, A. Shahroudy, G. Wang and D. Xu, *Skeleton-based action recognition using spatio-temporal LSTM network with trust gates*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2017.
- [CLLW15] M. Choo Chuah, W. Li, S. Lyu and L. Wen, *Category-blind human action recognition: a practical recognition system*, In IEEE International Conference on Computer Vision (ICCV), pp. 4444-4452, 2015.
- [CGILNPZ18] Christopher Clark, Matt Gardner, Mohit Iyyer, Kenton Lee, Mark Neumann, Matthew E Peters and Luke Zettlemoyer, *Deep contextualized word representations*, In Proceedings of NAACL-HLT, pages 2227-2237, 2018.
- [CJJNRSSW19] G. D. Clifford, R. Jeter, C. Josef, S. Nemati, M. A. Reyna, A. Sharma, S. P. Shashikumar and M. B. Westover, *Early prediction of sepsis from clinical data: the PhysioNet/Computing in Cardiology Challenge 2019*, Critical Care Medicine 2019, **48** (2), p.210-217, (2019).
- [CLNTZ16] T. Coleman, Y. Li, K. Nian, A. Tayal and H. Zhang, *Auto insurance fraud detection using unsupervised spectral ranking for anomaly*, Journal of Finance and Data Science, **2**(1):58-75, 2016.
- [Cut11] M. Cuturi, *Fast Global Alignment Kernels*, In Proceedings of the 28th international conference on machine learning (ICML-11), pp.929-936, 2011.
- [DR14] Snehal S Dahikar and Sandeep V Rode, *Agricultural crop yield prediction using artificial neural network approach*, International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering, **2**(1):683-686, 2014.
- [DHX03] S. Deng, Z. He and X. Xu, *Discovering cluster-based local outliers*, Pattern Recognition Letters, **24**(9-10):1641-1650, 2003.
- [DSS16] C. S. Deutschman, C. W. Seymour, M. Singer, et al., *The third international consensus definitions for sepsis and septic shock (Sepsis-3)*, JAMA 2016, **315**:801-810.
- [DFMS12] K. Dinuzzo, K. Fukumizu, K. Muandet and B. Schölkopf, *Learning from Distributions via Support Measure Machines*, In Advances in Neural Information Processing Systems, pages 10-18, 2012.
- [DWW15] Y. Du, W. Wang and L. Wang, *Hierarchical recurrent neural network for skeleton based action recognition*, In Computer Vision and Pattern Recognition (CVPR), pp. 1110-1118, 2015.

- [DG19] D. Dua and C. Graff, UCI machine learning repository, 2019.
- [DHKLW17] L. Y. Duan, P. Hu, A. C. Kot, J. Liu and G. Wang, *Global context-aware attention lstm networks for 3d action recognition*, In The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), vol. 7, pp. 43, 2017.
- [ELLY17] Stefano Ermon, Xiaocheng Li, David Lobell, Melvin Low and Jiaxuan You, *Deep gaussian process for crop yield prediction based on remote sensing data*, In Thirty-First AAAI Conference on Artificial Intelligence, 2017.
- [ESW10] F. Eyben, M. Wollmer and B. Schuller, *Opensmile: the munich versatile and fast open-source audio feature extractor*, Proceedings of the 18th ACM international conference on Multimedia. ACM, pp. 1459–1462, 2010.
- [FTZ13] J. Fan, X. Tian and Z. Zha, *Action recognition with novel high-level pose features*, In IEEE International Conference on Multimedia and Expo Workshops (ICMEW), pp. 1-6, 2016.
- [FLT17] H. Fang, C. Lu, Y. W. Tai and S. Xie, *Rmpe: Regional multi-person pose estimation*, In IEEE International Conference on Computer Vision (ICCV), vol. 2, 2017.
- [Faw03] Thomas Fawcett, *Problems in stochastic analysis: connections between rough paths and non-commutative harmonic analysis*, PhD Thesis, University of Oxford, 2003.
- [FLWZZ13] R. Fergus, Y. LeCun, L. Wan, M. Zeiler and S. Zhang, *Regularization of neural networks using dropconnect*, In International Conference on Machine Learning (ICML), pp. 1058-1066, 2013.
- [Fla15] Seth Flaxman, *Machine Learning in Space and Time*, Ph. D. thesis, Carnegie Mellon University, 2015.
- [FLSS17] Seth Flaxman, Ho Chung Leon Law, Dino Sejdinovic and Dougal Sutherland, *Bayesian Approaches to Distribution Regression*, Proceedings International Conference on Artificial Intelligence and Statistics 2018.
- [FKLLO21] J. Foster, P. Kidger, X. Li, T. Lyons and H. Oberhauser, *Neural SDEs as Infinite-Dimensional GANs*, International Conference on Machine Learning (ICML) 2021.
- [FKLM20] J. Foster, P. Kidger, T. Lyons and J. Morrill, *Neural Controlled Differential Equations for Irregular Time Series*, Conference on Neural Information Processing Systems (NeurIPS) 2020.
- [FKLMS21] J. Foster, P. Kidger, T. Lyons, J. Morrill and C. Salvi, *Neural Rough Differential Equations for Long Time Series*, In International Conference on Machine Learning (ICML), pp. 7829-7838, 2021.
- [FKLL21] James Foster, Patrick Kidger, Xuechen Li and Terry Lyons, *Efficient and Accurate Gradients for Neural SDE*, NeurIPS 2021.
- [FLO20] James Foster, Terry Lyons and Harald Oberhauser, *An optimal polynomial approximation of Brownian motion*, SIAM Journal on Numerical Analysis, vol. 58, no. 3, 2020.
- [FHK04] P. Franti, V. Hautamaki and I. Karkkainen, *Outlier detection using k-nearest neighbour graph*, Proceedings of the 17th International Conference on Pattern Recognition, ICPR 2004, vol. 3, pp.430-433, IEEE, 2004.
- [FV10] P. K. Friz and N. B. Victoir, *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*, Cambridge University Press, 2010.
- [FN93] K. -I. Funahashi and Y. Nakamura, *Approximation of dynamical systems by continuous time recurrent neural networks*, Neural Networks, vol. 6, no. 6, pp.801-806, 1993.
- [GLM19] J. Gallacher, T. J. Lyons and P. J. Moore, *Using path signatures to predict a diagnosis of Alzheimer's disease*, for the Alzheimer's Disease Neuroimaging Initiative (2019), PLoS ONE 14(9): e0222212. <https://doi.org/10.1371/journal.pone.0222212>

- [GJR18] J. Gatheral, T. Jaisson and M. Rosenbaum, *Volatility is Rough*, Quantitative Finance, **18**:6, 933-949, 2018.
- [Gér19] Aurélien Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, O'Reilly Media, Second Edition, ISBN 978-1-492-03264-9, 2019.
- [GLMS16] S. Ghosh, E. Laksana, L. P. Morency and S. Scherer, *Representation learning for speech emotion recognition*, Interspeech 2016, pp. 3603-3607, 2016.
- [GB18] Benjamin Graham and Jeremy Reizenstein, *The iisignature library: efficient calculation of iterated integral signatures and log signatures*, ACM Transactions on Mathematical Software 2020.
- [GPSS16] A. Gretton, B. Póczos, B. K. Sriperumbudur and Z. Szabó, *Learning Theory for Distribution Regression*, The Journal of Machine Learning Research, **17**(1):5272-5311, 2016.
- [GSSS07] A. Gretton, B. Schölkopf, A. Smola and L. Song, *A Hilbert Space Embedding for Distributions*, In International Conference on Algorithmic Learning Theory, pages 13-31, Springer, 2007.
- [Gur97] Kevin Gurney, *An Introduction to Neural Networks*, CRC Press, ISBN 978-1-85728-503-1, 1997
- [HNP09] Alon Halevy, Peter Norvig and Fernando Pereira, *The Unreasonable Effectiveness of Data*, Intelligent Systems, IEEE **24**(2):8 - 12, 2009.
- [HL10] Ben Hambly and Terry Lyons, *Uniqueness for the signature of a path of bounded variation and the reduced path group*, Annals of Mathematics, Volume **171**, pp.109-167, 2010.
- [HW07] Q. P. He and J. Wang, *Fault detection using the k-nearest neighbor rule for semiconductor manufacturing processes*, IEEE Trans. Semicond. Manuf., **20**, pp.345-354, 2007.
- [Hea12] Jeff Heaton, *Introduction to the Math of Neural Networks*, Heaton Research, 2012.
- [HHMR20] C. van der Heide, L. Hodgkinson, M. Mahoney and F. Roosta, *Stochastic Continuous Normalizing Flows*, Conference on Uncertainty in Artificial Intelligence (UAI) 2021.
- [Hil86] Terrell L. Hill, *An introduction to statistical thermodynamics*, Courier Corporation, 1986.
- [HKLMNS19] S. Howison, A. Kormilitzin, T. Lyons, J. Morrill, A. Nevado-Holgado and S. Swaminathan, *The signature-based model for early detection of sepsis from electronic health records in the intensive care unit*, In: 2019 Computing in Cardiology (CinC) (2019).
- [HKLMNS20] S. Howison, A. Kormilitzin, T. Lyons, J. Morrill, A. Nevado-Holgado and S. Swaminathan, *Utilization of the signature method to identify the early onset of sepsis from multivariate physiological time series in critical care monitoring*, Critical Care Medicine **48**(10), 976-981, 2020.
- [HLLW17] Y. Hou, C. Li, W. Li and P. Wang, *Joint distance maps Based action recognition with convolutional neural networks*, IEEE Signal Processing Letters, vol. **24**, no. 5, pp. 624-628, 2017.
- [HLZZ15] J. F. Hu, J. Lai, J. Zhang and W. S. Zheng, *Jointly learning heterogeneous features for rgb-d activity recognition*, In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5344-5352, 2015.
- [Hur51] H. Hurst, *The Long-Term Storage Capacity Reservoirs*, Transactions of the American Society of Civil Engineers, 1951.
- [JLNSY17] Lianwen Jin, Terry Lyons, Hao Ni, Cordelia Schmid and Weixin Yang, *Developing the Path Signature Methodology and its Application to Landmark-based Human Action Recognition*, In: Yin, G., Zariphopoulou, T. (eds) Stochastic Analysis, Filtering and Stochastic Optimization, Springer, Cham. 2022. https://doi.org/10.1007/978-3-030-98519-6_18

- [JLNY16] Lianwen Jin, Terry Lyons, Hao Ni and Weixin Yang, *Rotation-free online handwritten character recognition using dyadic path signature features, hanging normalization, and deep neural network*, 23rd International Conference on Pattern Recognition (ICPR), IEEE, pp. 4083–4088, 2016.
- [Joh73] G. Johansson, *Visual perception of biological motion and a model for its analysis*, Perception and psychophysics, vol **14**, no.2, pp.201-211, 1973.
- [JS00] A. K. Jones and R. S. Sielken, *Computer system intrusion detection: A survey*, Technical report, University of Virginia, 2000.
- [KN14] I. Kapsouras and N. Nikoladis, *Action recognition on motion capture data using a dynemes and forward differences representation*, Journal of Visual Communication and Image Representation, vol. **25**, no. 6, pp. 1432-1445, 2014.
- [KS20] Abbas Karamodin and Hassan Sarmadi, *A novel anomaly detection method based on adaptive Mahalanobis-squared distance and one-class kNN rule for structural health monitoring under environmental effects*, Mechanical Systems and Signal Processing, Volume **140**, June 2020.
- [KLMY21] Patrick Kidger, Terry Lyons, James Morrill and Lingyi Yang, *Neural Controlled Differential Equations for Online Prediction Tasks*, arXiv preprint, 2021. <https://arxiv.org/abs/2106.11028>
- [KL20] Patrick Kidger and Terry Lyons, *Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU*, International Conference on Learning Representations (ICLR) 2021.
- [KO19] Franz J. Király and Harald Oberhauser, *Kernels for sequentially ordered data*, Journal of Machine Learning Research, 2019.
- [KLMMSSSTVWWW11] Iwona Kloszewska, Simon Lovestone, Patrizia Mecocci, J-Sebastian Muehlboeck, Andrew Simmons, Hilkka Soininen, Christian Spenger, Magda Tsolaki, Bruno Vellas, Lars-Olof Wahlund, Michael W. Weiner and Eric Westman, *AddNeuroMed and ADNI: similar patterns of Alzheimer’s atrophy and automated MRI classification accuracy in Europe and North America*, Neuroimage **58**(3), 818-828, 2011.
- [KT01] D. Koller and S. Tong, *Support Vector Machine Active Learning with Applications to Text Classification*, Journal of Machine Learning Research, **2**, pp.45-66, 2001.
- [KPRSSZ17] Satwik Kottur, Barnabas Poczos, Siamak Ravanbakhsh, Russ R. Salakhutdinov, Alexander J. Smola and Manzil Zaheer, *Deep Sets*, In Advances in neural information processing systems, pages 3391–3401, 2017
- [LLLNO9] L. Lacasa, B. Luque, J. Luque and J. C. Nuno, *The visibility graph: A new method for estimating the Hurst exponent of fractional Brownian motion*, EPL (Europhysics Letters), vol.**86**, no.3, p.30001, 2009.
- [LP11] Christophe Ladroue and Anastasia Papavasiliou, *Parameter estimation for rough differential equations*, The Annals of Statistics, **39**(4):2047–2073, 2011.
- [LLSXXZZ16] C. Lan, Y. Li, L. Shen, X. Xie, J. Xing, W. Zeng and W. Zhu, *Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks*, In AAAI, vol. **2**, 2016.
- [LLSXZ17] C. Lan, J. Liu, S. Song, J. Xing and W. Zeng, *An end-to-end spatio-temporal attention model for human action recognition from skeleton data*, In AAAI, vol. **1**, no. 2, p. 7, 2017.
- [LT15] J. Lee and I. Tashev, *High-level feature representation using recurrent neural network for speech emotion recognition*, Sixteenth Annual Conference of the International Speech Communication Association, 2015.
- [Lee60] M. Lees, *The Goursat Problem*, Journal of the Society for Industrial Applied Mathematics, **8** (1960), pp. 518-530, 1960.

- [LLN16] Daniel Levin, Terry Lyons and Hao Ni, *Learning from the past, predicting the statistics for the future, learning an evolving system*, arXiv preprint, revised Mar 2016. <https://arxiv.org/abs/1309.0260>
- [LLNSW19] Maria Liakata, Terry Lyons, Alejo J. Nevado-Holgado, Hao Ni, Kate Saunders and Bo Wang, *A Path Signature Approach for Speech Emotion Recognition*, Interspeech 2019.
- [LLNSTWW20] Maria Liakata, Terry Lyons, Alejo J. Nevado-Holgado, Kate Saunders, Niall Taylor, Bo Wang and Yue Wu, *Learning to Detect Bipolar Disorder and Borderline Personality Disorder with Language and Speech in Non-Clinical Interviews*, Interspeech 2020.
- [LLSVWW21] Maria Liakata, Terry Lyons, Kate E. Saunders, Nemanja Vaci, Bo Wang and Yue Wu, *Modelling Paralinguistic Properties in Conversational Speech to Detect Bipolar Disorder and Borderline Personality Disorder*, ICASSP 2021.
- [LL11] Christian Litterer and Terry Lyons, *Cubature on Wiener space*, Stochastic Processes and Applications to Mathematical Finance, 197-218, 2011.
- [LTZ08] F. T. Liu, K. M. Ting and Z. H. Zhou, *Isolation forest*, In Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, pages 413-422, IEEE, 2008.
- [LTZ12] F. T. Liu, K. M. Ting and Z. H. Zhou, *Isolation-based anomaly detection*, ACM Transaction on Knowledge and Discovery from Data (TKDD), **6**(1):1-39, 2012.
- [LT18] G. Liu and F. Tao, *Advanced LSTM: A study about better time dependency modelling in emotion recognition*, ICASSP 2018, pp. 2906-2910, 2018.
- [LNSW16] J. Liu, T. T. Ng, A. Shahroudy and G. Wang, *NTURGB+D: A large scale dataset for 3D human activity analysis*, In Computer Vision and Pattern Recognition (CVPR), pp. 1010-1019, 2016.
- [LSWX16] Jun Liu, Amir Shahroudy, Gang Wang and Dong Xu, *Spatio-temporal LSTM with trust gates for 3D human action recognition*, In European Conference on Computer Vision (ECCV), pp. 816-833, 2016.
- [LPY19] Zhiyong Lu, Yifan Peng and Shankai Yan, *Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets*, In Proceedings of the 18th BioNLP Workshop and Shared Task, pages 58-65, 2019.
- [Lyo98] Terry Lyons, *Differential equations driven by rough signals*, Revista Matemática Iberoamericana, **14**(2): 215-310, 1998.
- [Lyo02] Terry Lyons, *ESig*. <https://esig.readthedocs.io/en/latest/>.
- [Lyo14] Terry Lyons, *Rough paths, Signatures and the modelling of functions on streams*, International Congress of Mathematicians (ICM) 2014.
- [LCL04] Terry Lyons, Michael J. Caruana and Thierry Lévy, *Differential Equations Driven By Rough Paths*, Springer, Ecole d'Été de Probabilités de Saint-Flour XXXIV, 2004.
- [LN11] Terry Lyons and Hao Ni, *Expected signature of two dimensional Brownian Motion up to the first exit time of the domain*, The Annals of Probability, Vol. **43**, No. 5, 2729-2762, 2015.
- [LY13] Terry Lyons and Danyu Yang, *The partial sum process of orthogonal expansions as geometric rough process with fourier series as an example - An improvement of Menshov-Rademacher theorem*, Journal of Functional Analysis **265**, no. 12, 3067-3103, 2013.
- [LV04] Terry Lyons and Nicolas Victoir, *Cubature on Wiener Space*, Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **460** (2004), no. 2041, 169-198, 2004.
- [Mah27] Prasanta Chandra Mahalanobis, *Analysis of race mixture in Bengal*, Journal and Proceedings of the Asiatic Society of Bengal, Vol. **23**, p.301-333, 1927.

- [Mah36] Prasanta Chandra Mahalanobis, *On the generalised distance in statistics*, Proceedings of the National Institute of Sciences of India, Vol. 2(1), p.49-55, 1936.
- [Mic10] Microsoft, *Kinect*, 2010. <https://www.microsoft.com/en-gb/p/azure-kinect-dk/8pp5vxmd9nhq?activetab=pivot%3aoverviewtab>
- [NCEZID16] Centers for Disease Control and Prevention: National Center for Emerging and Zoonotic Infectious Diseases (NCEZID), Division of Healthcare Quality Promotion (DHQP), 2016. <https://www.cdc.gov/sepsis/dataareports/index.html>
- [NV17] M. Neumann and N. T. Vu, *Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length and acted speech*, Proc. Interspeech 2017, pp. 1263-1267, 2017.
- [OT20] Harald Oberhauser and Csaba Toth, *Bayesian Learning from Sequential Data using Gaussian Processes with Signature Covariances*, ICML in press (2020).
- [RV14] V. B. Radhakrishnan and S. Vantigodi, *Action recognition from motion capture data using meta-cognitive rbf network classifier*, In International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), pp. 1-6, 2014.
- [RT19] M. Raginsky and B. Tzen, *Neural Stochastic Differential Equations: Deep Latent Gaussian Models in the Diffusion Limit*, arXiv preprint, 2019. <https://arxiv.org/abs/1905.09883>
- [Rei99] Linda E. Reichl, *A modern course in statistical physics*, 1999.
- [RSWX19] Kirk Roberts, Yuqi Si, Jingqi Wang and Hua Xu, *Enhancing clinical concept extraction with contextual embeddings*, Journal of the American Medical Informatics Association, Volume 2 (Short Papers), pages 464-468, 2019.
- [SS09] R. Sankar and N. I. Sapankevych, *Time Series Prediction Using Support Vector Machines: A Survey*, IEEE Computational Intelligence Magazine, 4, pp.24-38, 2009.
- [SW13] C. Schmid and H. Wang, *Action recognition with improved trajectories*, In IEEE International Conference on Computer Vision (ICCV), pp. 3551-3558, 2013.
- [SS18] B. Scholkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, Adaptive Computation and Machine Learning Series, 2018.
- [Sch89] Erwin Schrödinger, *Statistical thermodynamics*, Courier Corporation, 1989.
- [Sto48] M. H. Stone, *The Generalized Weierstrass Approximation Theorem*, Mathematics Magazine, 21 (4): 167-184, 1948.
- [SV79] D. W. Stroock and S. R. S. Varadhan, *Multidimensional Diffusion Processes*, Springer Berlin Heidelberg, 1979.
- [Vap98] V. Vapnik, *The Support Vector Method of Function Estimation*, In Nonlinear Modeling, Springer, pp.55-85, 1998.
- [WWY13] C. Wang, Y. Wang and A. L. Yuille, *An approach to pose-based action recognition*, In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 915-922, 2013.
- [Wei17] E. Weinan, *A Proposal on Machine Learning via Dynamical Systems*, Commun. Math. Stat., vol. 5, no. 1, pp.1-11, 2017.

UNIVERSITY OF OXFORD, RADCLIFFE OBSERVATORY, ANDREW WILES BUILDING, WOODSTOCK RD,
OXFORD, OX2 6GG, UK.

TL: tlyons@maths.ox.ac.uk

<https://www.maths.ox.ac.uk/people/terry.lyons>

AM: andrew.mcleod@maths.ox.ac.uk

<https://www.maths.ox.ac.uk/people/andrew.mcleod>