

Signature Methods in Machine Learning

Terry Lyons & Andrew D. McLeod

February 21, 2025

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 provide a “lossy compression” of the information contained within such a stream by filtering out the parameterisation noise. More concretely, suppose we have a time series with 3 channels and N samples. There are $1 + 3N + \frac{3N(3N+1)}{2} = 1 + \frac{9}{2}N(N+1)$ linearly independent quadratic polynomials defined on the time series. But the signature of this time series truncated to depth 2 only consists of $1 + 3 + 3^2 = 13$ components which is, in particular, independent of the number of samples N . However, whilst the independence of the number of samples N has removed an exponential amount of noise, the dependence on the number of channels to the power of the depth ensures that an exponential amount of information remains.

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.

Mathematical Subject Classification 2020: **60L10** (Primary); **93C15, 68Q32, 34F05** (Secondary).

Contents

1	Introduction	2
2	Background Material	5
2.1	Illustrative Toy Regression Problem	5
2.2	Incremental Streamed Data	6
2.3	Mathematical Framework	9

2.4	Simple Regression Approach	11
2.5	Streams of Increments as Paths	13
2.6	Signature Involvement	15
2.7	Tensor Algebra and Signature	17
2.8	Signature of Paths with Finite p-Variation	20
2.9	Rough Paths	24
3	Signature as a Feature Map	28
4	Controlled Differential Equations and the Log-ODE Method	32
5	Computing Signatures	34
6	Expected Signature	39
7	Truncation Order Selection	42
8	Signature Kernel	45
9	Neural CDEs and Neural RDEs	49
10	Speech Emotion Recognition	53
11	Health Applications	55
11.1	Bipolar and Borderline Personality Disorders	55
11.2	Alzheimer's Disease	58
11.3	Early Sepsis Detection	60
11.4	Information Extraction from Medical Prescriptions	63
12	Landmark-based Human Action Recognition	66
12.1	Path Disintegration and Transformations	67
12.2	Datasets	68
12.3	Method Implementation	69
12.4	Comparison with Specifically-Tailored Methods	71
12.5	Demo Notebook	72
13	Distribution Regression via the Expected Signature	73
14	Conformance and SigMahaKNN Method for Anomaly Detection	76
15	Randomised Signature	83

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.

A stream of sequential information can be interpreted as a path. Reparameterisation then gives an infinite dimensional group of symmetries. Typically symmetries are problematic for machine learning; different

representations of the same object introduce an additional fact that the model must learn. One route around this issue for paths is to work with the *unparameterised path*; that is, the equivalence class with respect to the group of reparameterisations (sometimes called a *shape* or a *contour*). Only the data itself is important and not the particular choice of representation; we only want to extract information provided by the data that is *independent* of the choice of how the data is viewed. For example, we may be interested in which character a sequence of pen strokes results in, but it is of no importance how quickly the character is drawn. A figure '3' is the same regardless of how quickly it is written.

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. The parameterisation noise is filtered out. 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 far 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 [Chevyrev and Kormilitzin, 2016]. 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 [Chevyrev and Kormilitzin, 2016] to cover more sophisticated uses of the signature transform in machine learning made after the appearance of [Chevyrev and Kormilitzin, 2016].

In the interest of readability we do not present this article as a proper sequel to [Chevyrev and Kormilitzin, 2016]. 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 [Chevyrev and Kormilitzin, 2016]. Subsection 2.1 provides a very basic introduction to regression via a simple example of monomial regression on the unit interval $[0, 1] \subset \mathbb{R}$. Subsection 2.2 presents a sensible and informative way to convert a data stream of values to a data stream of increments. In particular, no information is lost under this transform. Subsection 2.3 fixes the mathematical framework for the regression problem that we will consider throughout this article.

Subsection 2.4 introduces some basic strategies to tackling the regression problem introduced in Subsection 2.3. The viewpoint adopted in Subsection 2.4 is that of a beginner who has no prior experience with machine learning. The machine learning material included in Section 2.4 is only the essentials we later require. The reader seeking a more comprehensive introduction to machine learning is directed to the book [Géron, 2022]. The book [Géron, 2022] covers the fundamentals of *machine learning*, *classification*, *regression*, *neural networks* and *deep learning* with exercises and coded examples. The more recent book [Raschka et al., 2022] covers a broadly similar collection of topics whilst also providing coded examples. A major difference between [Géron, 2022] and [Raschka et al., 2022] is that the coded examples in [Géron, 2022] use *TensorFlow*, whilst the coded examples in [Raschka et al., 2022] use *PyTorch*. In addition, both [Géron, 2022] and [Raschka et al., 2022] include introductions to more recent developments within the field of data science including, for example, *large language models*, *attention*, and *transformers*. A more theoretical perspective of the field can be found in [Hastie et al., 2017], for example.

Subsection 2.5 explains the natural correspondence between a stream of increments and a continuous path. Additionally, Subsection 2.5 presents some simple augmentations that can be applied to streamed data. Only augmentations used in later applications are covered.

Subsection 2.6 informally introduces the signature transform as an attractive method of summarising a stream of increments via summarising the corresponding continuous path (cf. Subsection 2.5). After informally introducing the signature as a method to summarise a path in Section 2.6 that captures all the relevant

information, the mathematical framework required to rigorously talk about signatures is covered in Section 2.7. Keeping in mind our adopted machine learning perspective, signature properties and related theorems are all stated first for continuous paths of bounded p -variation for $1 \leq p < 2$ in Subsection 2.8. For the majority of this article we stay within the setting of considering signatures of continuous paths of bounded 1-variation; and so the main theoretical requirements are the results presented in Subsection 2.8.

However, in places we necessarily require the more involved theory of *rough paths*. Indeed, it is well known that the path $t \mapsto B_t$ for a standard Brownian motion B_t only has finite p -variation for $p > 2$. Thus the theory of rough paths is required to make sense of the signature of a path arising from the evolution of Brownian motion. Consequently we provide a terse presentation of the aspects of rough path theory that we will require in Subsection 2.9.

As a pseudo-sequel to [Chevyrev and Kormilitzin, 2016], we assume the reader has had some exposure to signatures. Our discussion of signatures assumes a greater mathematical maturity than [Chevyrev and Kormilitzin, 2016] and, in contrast to the example-heavy introduction given in [Chevyrev and Kormilitzin, 2016], 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 Subsections 2.7, 2.8 and 2.9. A wider range of properties of path signatures can be found in the survey article [Lyons, 2014], whilst the reader seeking a comprehensive introduction to the theory is directed to the lecture notes [Lyons et al., 2007].

The entirety of Section 2 is included with the aim of providing a dictionary for terminology and notation used throughout the later sections. The presentation of this material is not designed to provide an in-depth coverage; only the specific elements required to understand the later applications are covered. The reader already comfortable with *Regression*, *Feature Maps*, *Kernels*, *Rough Paths*, and the mathematical theory of *Signatures* may safely skim Subsections 2.1 – 2.9.

The material covered ceases over-lapping with the content of [Chevyrev and Kormilitzin, 2016] in Section 3. For the readers ease of navigation, we briefly summarise the content included in the remaining sections.

- Section 3 provides both a basic outline for the use of the signature transform as a feature map, and 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 Subsections 2.7, 2.8 and 2.9.
- Section 5 provides an introduction to the *RoughPy* [Morley and Lyons, 2024] python package 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 [Chevyrev and Oberhauser, 2022] for example. We additionally outline how this theory offers an approach to distribution regression following [Lemercier et al., 2021].
- Section 7 covers the work [Kidger et al., 2019] proposing how the signature transform may be incorporated into a neural network. A particular consequence is that the appropriate depth to truncate the signature to should be treated as a trainable parameter, 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 [Salvi et al., 2021] 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 [Kidger et al., 2020] and of *Neural Rough Differential Equations* (Neural RDEs) made in [Morrill et al., 2021b]. 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 [Wang et al., 2019] 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 [Perez Arribas et al., 2018, Wang et al., 2020], diagnosing Alzheimer’s disease [Moore et al., 2019], information extraction from medical prescriptions [Biyong et al., 2020], and early detection of Sepsis [Morrill et al., 2019, Morrill et al., 2020].
- Section 12 covers the work of [Yang et al., 2022] 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 [Yang et al., 2022]).
- Section 13 covers the work of [Lemercier et al., 2021] 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 [Shao et al., 2020].
- Section 15 covers *randomised signatures* as developed in the works [Cuchiero et al., 2021b, Akyildirim et al., 2022], and illustrates their use in [Akyildirim et al., 2022] for the detection of market manipulation attempts from financial data.

Whilst a wide range of applications are covered within this article, it remains a far from comprehensive presentation of the uses of path signature techniques within machine learning. Indeed the list of recent works *not* discussed within this article includes uses of signature techniques in finance [Cuchiero et al., 2023a, Neufeld and Schmock, 2022, Bank et al., 2023, Cuchiero and Möller, 2023], optimal stopping problems [Bayer et al., 2023c, Bayer et al., 2023d], stochastic differential equations [Cuchiero et al., 2023c], universal approximation theory [Cuchiero et al., 2022, Varzaneh et al., 2022, Cuchiero et al., 2023b], neural network stability [Bayer et al., 2023b], generalized Magnus expansion [Friz et al., 2022], functional Taylor expansions [Dupire and Tissot-Daguette, 2022, Dupire and Tissot-Daguette, 2025], economic nowcasting [Cohen et al., 2023], and longitudinal language modelling [Tseriotou et al., 2023b, Tseriotou et al., 2023a]. Moreover this list of works not included in this article is far from exhaustive. Including every topic within a single article is infeasible. But hopefully, after working their way through this article, the reader will be suitably equipped to read and understand other works considering the use of path signature techniques within a machine learning context.

Acknowledgements: This work was supported by the DataSig Program under the EPSRC grant ES/S026347/1, the Alan Turing Institute under the EPSRC grant EP/N510129/1, the Mathematical Foundations of Intelligence: An “Erlangen Programme” for AI under the EPSRC Grant EP/Y028872/1, the Data Centric Engineering Programme (under Lloyd’s Register Foundation grant G0095), the Defence and Security Programme (funded by the UK Government) and the Hong Kong Innovation and Technology Commission (InnoHK Project CIMDA). The example notebooks available via the DataSig website were created by Peter Foster with various members of the DataSig team. This work was funded by the Defence and Security Programme (funded by the UK Government). For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission. We thank the anonymous referee for both their careful reading of and valuable constructive feedback on an earlier version.

2. Background Material

2.1. Illustrative Toy Regression Problem

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. The following simple example of regression involving the consideration of the monomials on the unit interval $[0, 1] \subset \mathbb{R}$ exhibits some of the theoretical considerations that are central to the use of signatures within machine learning.

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 $N \in \mathbb{N}$. 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 rarely-used, yet still insightful, 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. Degeneracy in the equations can also be handled via regularisation techniques including, for example, *Ridge regression*, *LASSO regression*, and *dropout*. An in-depth discussion of such approaches may be found in [Géron, 2022] or [Raschka et al., 2022], for example.

The success of this approach to approximating ρ is fundamentally reliant on the linear span of the set $\{\phi_k \mid k \in \mathbb{N}\}$ 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 [Stone, 1948], for example) provides the required density.

2.2. Incremental Streamed Data

A fundamental idea central to data science, illustrated by the simple monomial example in Subsection 2.1, is that if we are able to represent data as vectors in some Euclidean space, then we can approximate continuous functions on compact domains by polynomials and learn the polynomial from the data. From a theoretical perspective this is valid; the *Stone–Weierstrass theorem* [Stone, 1948] tells us that, on compact subsets, polynomials are a dense subset of the continuous functions.

In practice this approach can work well when the data consists of time-series with moderate numbers of both channels and samples. Examples of tasks on which polynomial regression performs well include predicting the median house price for a district using 10 recorded attributes per district (see Chapter 2 in [Géron, 2022]), predicting the *Traction Energy Consumption* of an urban car journey based on knowledge of the time-of-day at which the journey is made [Kreichbergs and Grislis, 2019], and predicting the number of deaths resulting from COVID-19 in US states based on four independent parameters [Singh and Bawa, 2022].

But this approach runs into problems for tasks in which the order of events is important. In order to consider a toy problem illustrating one particular issue, let $N \in \mathbb{Z}_{\geq 1}$ and suppose we have a finite collection of time series, with each time series consisting of 2 channels and N samples. Suppose that each time series \mathbf{x} has the following structure. If $\mathbf{x} = \{(x_{1,i}, x_{2,i})\}_{i=1}^N$ then the following properties are true.

- For every $i \in \{1, \dots, N\}$ we have $x_{1,i}, x_{2,i} \in \{0, 1\}$.
- For $j \in \{1, 2\}$ if $i \in \{1, \dots, N\}$ and $x_{j,i} = 1$, then for every $k \in \{i, \dots, N\}$ we have $x_{j,k} = 1$.
- At least one of $x_{1,N}$ and $x_{2,N}$ is equal to 1.

Consider the task of determining which channel is the first to change from 0 to 1.

Each time series \mathbf{x} can be viewed as an element in \mathbb{R}^{2N} . We could try to learn a quadratic polynomial that can determine which channel is the first to change from 0 to 1. However, ignoring any considerations regarding whether a quadratic polynomial capable of doing this even exists, we observe that there are $1 +$

$3N + 2N^2$ linearly independent quadratic polynomials determined on \mathbf{x} . Consequently we must consider linear combinations of $1 + 3N + 2N^2$ elements in order to learn our desired quadratic polynomial. This quickly becomes intractable as the number of samples increases.

A large number of samples is reasonable. The sampling frequency needs to be high enough to separate the increases of the channels. If, for example, the time series correspond to a record of a patients heart rate and body temperature, with the change from 0 to 1 represents an increase in the heart rate or the body temperature respectively, then it may be the case that separating which one increases first could required taking a sample each second. This would result in $N = 3600$ samples over each hour, leading to the consideration of 25930801 linearly independent quadratic polynomials.

As we will see later (cf. Section 5), the signature provides a summary of time series that efficiently captures the order of events. In the toy problem above, the signature truncated to depth 2 can determine which channel changes from 0 to 1 first. This involves only 7 components rather than the $1 + 3N + 2N^2$ components of a linear combination of a basis of quadratic polynomials. The key advantage is that the number of signature components required is *independent* of the number of samples (though it does still grow exponentially with the number of channels).

Using the signature to summarise a time-series requires associating the time-series with a continuous path. A commonly used method is making a choice of rule for continuously interpolating between the entries of the time-series; examples include linear interpolation, rectilinear interpolation, or cubic spline interpolation. A weakness of this approach is that the choice of interpolation is typically ad-hoc; there is no given natural best choice. However, if the time-series is a stream of *increments*, then the concatenation of the entries provides a sensible route to determining a path with which the time-series can be associated (cf. Section 2.5). This natural association makes *incremental* streamed data (i.e. streams whose entries are all incremental changes) more desirable for use with path signature techniques.

The incremental change in a channels values is often vital information. For example, suppose we have a stream containing the recorded temperatures of a hospital patient. Observing the value 38°C alone is insufficient to conclude whether the patient is doing well or not. Knowledge of this values relation the the previous value is required. If the temperature was previously 39°C then the new value of 38°C suggests a positive trajectory for the patient; if the previous value of the temperature however was 37°C , then the new value 38°C suggests a more concerning trajectory for the patient, indicating that some form of medical intervention may be required.

Moreover, there are situations in which *only* the incremental changes are relevant, i.e. when the underlying actual values are not relevant. For example, doctors may be interested in knowing that a patient has taken 2 paracetamol tablets since yesterday, but they will not be interested in knowing the precise number of tablets the patient has taken over their entire lifetime. Similarly, in a financial context, the amount of volatility experienced over the past hour can be informative; the total amount of volatility experienced since the beginning-of-time is not. When planning a journey, the number of buses from a particular stop in the next hour could be helpful; but knowing how many buses have ever departed from that stop is likely not helpful. An interesting facet of these examples is that the irrelevancy of the actual value is, at least in part, due to its lack of availability.

However, it is important to realise that there are situations in which the incremental changes alone are insufficient. Indeed in the hospital patient temperature records example above, observing the incremental change of $+0.5^\circ\text{C}$ is insufficient to conclude whether the patient is doing well. If the last value before this change was 35°C then this increase is a good sign. But if the last value before this change was 38°C then this increase is not a good sign, indicating that some form of medical intervention may be required.

An important takeaway from this example is that knowledge of both a quantities actual values *and* the incremental changes of the quantity are required to determine the context of the information provided by a stream. Consequently, with the aim of using path signature techniques to summarise streams, we seek a way to record/store *any* given stream as a stream of increments that encodes both the incremental changes of the values and the actual values themselves. A definition of a stream can be found in Subsection 2.3 (cf. (2.4)).

Before turning our attention to achieving this, we first observe a third aspect of a data stream that is desirable to capture. Namely, whether each entry of the stream is the result of a new measurement being taken, or if it is the result of the previous value being carried across without a new measurement being taken. The importance of this distinction can be again seen via the example of a hospital patients recorded temperatures. Suppose the same temperature of 36.5°C is recorded in successive entries. Then whether or not the second

entry is a new measurement provides information. If it is not a new measurement, then it indicates that the patients condition is deemed stable enough to not require a new check of their temperature. But if it is a new measurement, then it indicates that the patients condition was deemed concerning enough, for some unknown reason, to warrant a new check of their temperature.

We seek a method of transforming any stream into a stream of increments that satisfies the following properties.

- No information is lost; any information present in the original stream can be recovered from the resulting stream of increments.
- Both the incremental changes and the actual values of each channel of the original stream are encoded in the resulting stream of increments.
- If available, whether or not each entry of each channel of the original stream is the result of a new measurement or not is encoded in the resulting stream of increments.
- The transform is injective; distinct streams should be transformed to distinct streams of increments.

We now informally discuss a transformation satisfying these properties in the simple setting of having a stream $\mathbf{x} = (x_1, \dots, x_n)$ of real numbers; that is, for every $i \in \{1, \dots, n\}$ we have $x_i \in \mathbb{R}$. The first step is to mimic a “pen-on, pen-off” style transformation and introduce a second channel recording a 1 for each entry that is the result of a new measurement, and a 0 for each entry that is not the result of a new measurement. This results in a new stream $\mathbf{x}' = ((x_1, a_1), \dots, (x_n, a_n))$ such that, for each $i \in \{1, \dots, n\}$, $x_i \in \mathbb{R}$ and $a_i \in \{0, 1\}$. If it is unknown whether x_i is the result of a new measurement or not we take $a_i := 1$; by default, we assume that each recorded value is the result of a new measurement unless explicitly told otherwise.

With an eventual aim of converting the stream \mathbf{x}' into a stream of incremental changes, we first encode information about the values x_1, \dots, x_n in a manner that is invariant under translation. We do this by explicitly encoding both the first value x_1 and the last value x_n via “invisibility-reset” type augmentations. To be precise, we replace the stream \mathbf{x}' by the new stream $\mathbf{x}' = ((0, 0), (x_1, 0), (x_1, a_1), \dots, (x_n, a_n), (x_n, 0), (0, 0))$. The “invisibility-reset” aspect to encode the first value x_1 and the last value x_n is done via a “Lead-Lag” type augmentation in which one channel updates, and then there is a definite delay until the other channel updates. In this case, at the start we first update 0 to x_1 and then update 0 to a_1 after a delay; whilst at the end we update a_n to 0, and then update x_n to 0 after a delay.

This “Lead-Lag” delayed updating (cf. Subsection 2.5) is a simple example of a sliding variable giving information at a point about what is happening on a neighbourhood of the point. Unfortunately it accidentally fixes a parameterisation of the data set via the explicit choice of the number of delay steps. It would be possible to design time-invariant contextual variables to provide information at a point about what is happening on a neighbourhood of the point. For example, taking the signature of a path over previous interval started at the time the path *last* entered the closed unit ball centred at the current location would give a more path intrinsic method.

Choosing to explicitly encode the first and last values is natural from a modelling perspective. It is these values that are particularly informative for a wide variety of tasks. For example, if our aim is to predict the subsequent values of a quantity (i.e. predict the next values of stock prices, patients body temperature etc) then the last known value of the quantity will be of particular significance. But if the task is classify the outcome of a quantities evolution (i.e. medical diagnosis based on medical records, whether an urban areas infrastructure will be sufficient to deal with demand etc) then the first value will be of particular significance. In these cases it is desirable to understand from the first recorded measurement what the likely outcome will be.

Moreover, the first and last entries of the stream are distinct from every other entry in that they only have one neighbouring entry. Every other entry has both a proceeding entry and a subsequent entry with which it can be compared. The first entry has no proceeding entry, and the last entry has no subsequent entry. Augmenting the stream by adding in additional entries to proceed the first entry and follow the last entry is a sensible way to make the first and last entries more uniform with respect to the other entries.

We choose to use the value 0 to both proceed x_1 and succeed x_n . As mentioned previously, one benefit of this is that the value x_1 and x_n will be explicitly encoded in a translation invariant manner. In particular, they

will remain encoded after the stream is replaced by its stream of incremental changes. Another benefit is that comparing x_1 and x_n with 0 makes sense for all data types; it represents comparison with the default setting of not having a measurement. Whilst in certain contexts it is arguably more sensible to compare these values with some notion of average (i.e. compare body temperature with average human body temperature value), there are settings in which such an average value does not make sense. For example, if the stream represents a series of x-rays of a patient, it is unclear what meaning should be given to the average x-ray. But the notion of not having an x-ray does make sense.

Finally, the stream \mathbf{x}' is replaced by the stream of its incremental differences. That is, we set

$$\mathbf{x}' := \left((x_1, 0), (0, a_1), (x_2 - x_1, a_2 - a_1), \dots, (x_n - x_{n-1}, a_n - a_{n-1}), (0, -a_n), (-x_n, 0) \right). \quad (2.3)$$

This transformation takes a stream with one channel to a stream with two channels. For a general stream with d channels, we will apply the transformation outlined above to each channel. Hence overall it will result in the stream with d channels being transformed to a stream with $2d$ channels (cf. Subsection 2.3).

We end this subsection with a brief discussion of some possible redundancies in this transformation. The first is that if the value of the entries of the stream are known to be of no importance, then avoiding the invisibility-reset aspect of the transformation is a more cost-effective option. The second is that we only really require implementing the invisibility-reset aspect to one of x_1 or x_n . After it has been implemented to one of these values, the other value is then recoverable from the resulting stream of increments. Our choice to include this aspect for both the first and last entries is made for two main reasons. Firstly, it emphasises the extra importance placed on the first and last values. Secondly, it makes our transformation more symmetrical, in particular ensuring that the concatenation of the resulting increments forms a closed path based at the origin.

2.3. Mathematical Framework

In this subsection we introduce the mathematical framework that we work in throughout this article. Given an arbitrary set \mathcal{M} we define the collection of *streams* in \mathcal{M} to be the set

$$\mathcal{S}(\mathcal{M}) := \bigcup_{k=1}^{\infty} \mathcal{M}^k = \bigcup_{k=1}^{\infty} \{(x_1, \dots, x_k) \mid x_1, \dots, x_k \in \mathcal{M}\}. \quad (2.4)$$

Throughout this article we will consider streams in the setting that the set \mathcal{M} is a real Banach space (complete normed linear vector space) that we denote by W . Moreover, we will assume that this Banach space W is finite dimensional with dimension $d \in \mathbb{Z}_{\geq 1}$.

We consider the following regression problem. Suppose that $M \in \mathbb{Z}_{\geq 1}$ is finite and we have a collection of pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^M$ where, for each $i \in \{1, \dots, M\}$, $\mathbf{x}_i \in \mathcal{S}(W)$ and $y_i \in \mathbb{R}$. The underlying idea is that for each $i \in \{1, \dots, M\}$ the value $y_i \in \mathbb{R}$ represents the underlying systems response to the stream \mathbf{x}_i . Suppose we have a subset $\mathcal{M} \subset \mathcal{S}(W)$ such that $\{\mathbf{x}_1, \dots, \mathbf{x}_M\} \subset \mathcal{M}$. Then we can consider the task of transitioning from the collection of pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^M$ to a continuous function $f : \mathcal{M} \rightarrow \mathbb{R}$ that accurately reflects the known responses of the system to the inputs $\mathbf{x}_1, \dots, \mathbf{x}_M$ and is capable of predicting the response to new inputs.

Making sense of a continuous function $f : \mathcal{M} \rightarrow \mathbb{R}$ requires a topology on \mathcal{M} . One particular way to equip \mathcal{M} with a topology is to choose a topology for $\mathcal{S}(W)$ and then have \mathcal{M} inherit this topology as a subset. There are, of course, several different ways this could be done. As an example, we could define a norm on $\mathcal{S}(W)$ as follows. If $k \in \mathbb{Z}_{\geq 1}$ and $\mathbf{x} = (x_1, \dots, x_k) \in \mathcal{S}(W)$ then its norm is

$$\|\mathbf{x}\|_{\mathcal{S}(W)} := \sum_{j=1}^k \|x_j\|_W. \quad (2.5)$$

This choice corresponds to using the $l^1(W)$ norm on the space of streams $\mathcal{S}(W)$, observing that any element $\mathbf{x} \in \mathcal{S}(W)$ can be viewed as an element in $l^1(W)$ whose entries are all eventually 0. Making an appropriate choice of norm for $\mathcal{S}(W)$ is a modelling problem. There is no choice that will always be the most sensible.

Which choice of norm is most suitable will depend on the particular problem.

We approach this regression problem by first transforming each of the streams $\mathbf{x}_1, \dots, \mathbf{x}_M \in \mathcal{S}(W)$ to a stream of increments using the transformation informally introduced at the end of Subsection 2.2. We first formalise this transformation. Consider a basis $e_1, \dots, e_d \in W$ for W . Suppose $k \in \mathbb{Z}_{\geq 1}$ and let $\mathbf{x} = (x_1, \dots, x_k) \in \mathcal{S}(W)$ be a stream such that, for each $j \in \{1, \dots, k\}$, we have

$$x_j = \sum_{l=1}^d x_{j,l} e_l \quad (2.6)$$

for real coefficients $x_{j,1}, \dots, x_{j,d} \in \mathbb{R}$. If we let the j^{th} row correspond to the coefficient of the basis element e_j , we can represent the stream \mathbf{x} as

$$\mathbf{x} = \left\{ \begin{pmatrix} x_{1,1} \\ \vdots \\ x_{1,d} \end{pmatrix}, \dots, \begin{pmatrix} x_{k,1} \\ \vdots \\ x_{k,d} \end{pmatrix} \right\}. \quad (2.7)$$

Then we apply the transformation introduced at the end of Subsection 2.2 to each channel in (2.7). That is, we transform the stream \mathbf{x} as represented in (2.7) to the new stream

$$\mathbf{x}' := \left\{ \begin{pmatrix} x_{1,1} \\ 0 \\ \vdots \\ x_{1,d} \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -a_{1,1} \\ \vdots \\ 0 \\ -a_{1,d} \end{pmatrix}, \begin{pmatrix} x_{2,1} - x_{1,1} \\ a_{2,1} - a_{1,1} \\ \vdots \\ x_{2,d} - x_{1,d} \\ a_{2,d} - a_{1,d} \end{pmatrix}, \dots, \begin{pmatrix} x_{k,1} - x_{k-1,1} \\ a_{k,1} - a_{k-1,1} \\ \vdots \\ x_{k,d} - x_{k-1,d} \\ a_{k,d} - a_{k-1,d} \end{pmatrix}, \begin{pmatrix} 0 \\ -a_{k,1} \\ \vdots \\ 0 \\ -a_{k,d} \end{pmatrix}, \begin{pmatrix} -x_{k,1} \\ 0 \\ \vdots \\ -x_{k,d} \\ 0 \end{pmatrix} \right\}. \quad (2.8)$$

For every $i \in \{1, \dots, k\}$ and $j \in \{1, \dots, d\}$ the real number $a_{i,j} \in \{0, 1\}$ records whether the recorded value $x_{i,j} \in \mathbb{R}$ is the result of a new measurement or not. We have $a_{i,j} = 1$ if $x_{i,j}$ is the result of a new measurement, and $a_{i,j} = 0$ otherwise. If this information is unknown or not available, then the default value is to take $a_{i,j} = 1$. That is, unless told otherwise we assume that each entry $x_{i,j}$ is the result of a new measurement.

The values of the first and last entries of the original stream \mathbf{x} are encoded in the first and last entries of the new stream \mathbf{x}' respectively. The particular importance of the first and last values of a stream are emphasised by this explicit encoding. In fact, the value of any entry of the original stream \mathbf{x} is recoverable from the new stream \mathbf{x}' . Moreover, if two streams \mathbf{x}_1 and \mathbf{x}_2 are distinct, then the new streams \mathbf{x}'_1 and \mathbf{x}'_2 , resulting from the application of this transformation to \mathbf{x}_1 and \mathbf{x}_2 respectively, remain distinct. No information is lost under this transformation. Consequently this transformation of a stream to a stream of increments satisfies the wishlist of properties discussed in Subsection 2.2.

It is seen from (2.8) that the stream \mathbf{x}' is an element in $\mathcal{S}(W \times \mathbb{R}^d)$. Let $V := W \times \mathbb{R}^d$ and replace each of the streams $\mathbf{x}_1, \dots, \mathbf{x}_M \in \mathcal{S}(W)$ by the stream of increments in $\mathcal{S}(V)$ resulting from applying the transformation determined by (2.8). Making a slight abuse of notation, we continue to denote the resulting streams of increments in $\mathcal{S}(V)$ by $\mathbf{x}_1, \dots, \mathbf{x}_M \in \mathcal{S}(V)$. Replace each stream \mathbf{x} in the subset $\mathcal{M} \subset \mathcal{S}(V)$ by the stream resulting from applying the transformation determined in (2.8) to \mathbf{x} . Again we abuse notation and continue to denote the resulting subset of $\mathcal{S}(V)$ by \mathcal{M} . That is, we now have $\mathcal{M} \subset \mathcal{S}(V)$.

We now consider the following regression problem. We have a finite collection of pairs $\Omega = \{(\mathbf{x}_i, y_i)\}_{i=1}^M$ where, for each $i \in \{1, \dots, M\}$, \mathbf{x}_i is a stream of increments in V and $y_i \in \mathbb{R}$ represents the underlying systems response to the stream \mathbf{x}_i . We adopt the convention that Ω is a *labelled dataset*. We denote by $\Omega_{\mathcal{S}(V)}$ the projection of Ω to $\mathcal{S}(V)$, that is $\Omega_{\mathcal{S}(V)} := \{\mathbf{x}_i \mid i \in \{1, \dots, M\}\} \subset \mathcal{S}(V)$. We adopt the convention that $\Omega_{\mathcal{S}(V)}$ is a *dataset*. Observe that $\Omega_{\mathcal{S}(V)} \subset \mathcal{M} \subset \mathcal{S}(V)$. We consider the task of transitioning from the collection of pairs Ω to a continuous function $f : \mathcal{M} \rightarrow \mathbb{R}$ that accurately reflects the known response of the system to the inputs $\mathbf{x}_1, \dots, \mathbf{x}_M$, and is capable of predicting the response to new inputs $\mathbf{x} \in \mathcal{S}(V)$ that do not correspond to an instance within $\Omega_{\mathcal{S}(V)}$.

Once again making sense of a continuous function $f : \mathcal{M} \rightarrow \mathbb{R}$ requires a choice of topology on \mathcal{M} .

This can again be achieved by equipping $\mathcal{S}(V)$ with a topology and letting \mathcal{M} inherit the subspace topology. There are, of course, numerous possible ways to equip $\mathcal{S}(V)$ with a topology. As an example, we could first choose a norm N for \mathbb{R}^d . Then we could define a norm $\|\cdot\|_V$ on $V = W \times \mathbb{R}^d$ by defining $\|v\|_V := \|P_W(v)\|_W + N(P_d(v))$ for $v \in V$ where $P_W : V \rightarrow W$ and $P_d : V \rightarrow \mathbb{R}^d$ denote the respective projection maps. Then we could subsequently define a norm on $\mathcal{S}(V)$ as done for $\mathcal{S}(W)$ in (2.5). That is, if $k \in \mathbb{Z}_{\geq 1}$ and $\mathbf{x} = (x_1, \dots, x_k) \in \mathcal{S}(V)$ then its norm is

$$\|\mathbf{x}\|_{\mathcal{S}(V)} := \sum_{j=1}^k \|x_j\|_V. \quad (2.9)$$

Making an appropriate choice of norm for $\mathcal{S}(V)$ is again a modelling problem. There is no choice that will always be the most sensible. Which choice of norm is most suitable will depend on the particular problem.

In this formulation, we implicitly assume that the dataset $\Omega_{\mathcal{S}(V)}$ gives a “good approximation” (in some sense) of the subset \mathcal{M} . That is, we assume it is reasonable to try and learn the systems response to inputs in \mathcal{M} using only its responses to inputs in $\Omega_{\mathcal{S}(V)}$. 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 \mathcal{M} bearing little resemblance to *any* element in $\Omega_{\mathcal{S}(V)}$. 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 from 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 \mathcal{M} is somehow similar to an element in $\Omega_{\mathcal{S}(V)}$. 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 $\mathcal{M} \subset \mathcal{S}(V)$ is a given subset and that our dataset $\Omega_{\mathcal{S}(V)}$ provides a “sufficiently good” approximation of \mathcal{M} , without concerning ourselves with the precise details meant by “sufficiently good”.

2.4. Simple Regression Approach

We illustrate a basic approach to the regression problem outlined in Subsection 2.3, and highlight some of the key theoretical properties required for this approach to be successful. We will later present an approach to this regression task through the use of signatures, and validate that the required theoretical properties are satisfied (cf. Subsection 2.6 and Section 3). For examples of working regression into machine learning pipelines see [Géron, 2022] or [Raschka et al., 2022].

A core idea in many schemes for learning a function f from a collection of pairs Ω is to mimic the use of the monomials in the basic example presented in Subsection 2.1. That is, first identify basic functions or features that are readily evaluated at each point in $\Omega_{\mathcal{S}(V)}$. This role is played by the monomials on the unit interval $[0, 1] \subset \mathbb{R}$ in Subsection 2.1. Once a collection of basic functions has been identified, we try to express the observed function as a linear combination of these basic functions.

One approach to finding a suitable collection of feature functions is to identify a *universal feature map* Ψ . That is, to find an embedding $\Psi : \mathcal{M} \rightarrow W$ for a Banach space W for which

$$\overline{\Psi(\mathcal{M})}^* = C^0(\Psi(V)) \quad (2.10)$$

where the closure is taken with respect to the uniform C^0 -norm on W . That is, we take the closure of $\Psi(\mathcal{M})^*$ in $C^0(W; \mathbb{R})$ where $C^0(W; \mathbb{R})$ is equipped with the norm $\|A\|_{C^0(W; \mathbb{R})} := \sup_{w \in W} \{|A(w)|\}$.

A universal feature map Ψ provides a representation of the dataset $\Omega_{\mathcal{S}(V)}$ as a subset $\Omega_W := \Psi(\Omega_{\mathcal{S}(V)})$ of a Banach space W in which continuous functions can be approximated by linear functions. That is, assuming (2.10) is valid, if $F \in C^0(\Psi(\mathcal{M}))$ then given any $\varepsilon > 0$ there exists a linear functional $\sigma : \Omega_W \rightarrow \mathbb{R}$ such

that for every $w \in \Omega_W$ we have $|F(w) - \sigma(w)| \leq \varepsilon$. This would allow us to use linear functions as our feature functions. Before mimicking the approach of approximating a continuous function on the interval $[0, 1]$ described in Subsection 2.1, we need to understand how to evaluate $\phi(\mathbf{x}_i)$ for elements $\phi \in \Psi(\mathcal{M})^*$ and $i \in \mathcal{I}$.

We seek a suitable universal embedding of the set \mathcal{M} into some Banach space W . If we simultaneously found an embedding of \mathcal{M} into the dual space W^* , then we would have many of the properties of what the machine learning literature refers to as a *kernel*. To be more precise, a pair of embeddings

$$\Psi : \mathcal{M} \rightarrow W \quad \text{and} \quad \Phi : \mathcal{M} \rightarrow W^* \quad (2.11)$$

would allow us to define a function $K : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$ by setting

$$K(x, y) := \langle \Phi(x), \Psi(y) \rangle \quad (2.12)$$

where $\langle \cdot, \cdot \rangle : W^* \times W \rightarrow \mathbb{R}$ denotes the natural dual-pairing map. That is, given $\phi \in W^*$ and $v \in W$ we define $\langle \phi, v \rangle := \phi(v)$.

If W is a Hilbert space, in which case its dual W^* is isometric to itself and the natural dual-pairing map is simply the inner product on W , then the function K defined in (2.12) is a kernel. In this setting the kernel K generates what is called a *Reproducing Kernel Hilbert Space* (RKHS), denoted by \mathcal{H}_K , that is a subset of the continuous functions $\mathcal{M} \rightarrow \mathbb{R}$, i.e. $\mathcal{H}_K \subset C^0(\mathcal{M}; \mathbb{R})$. Originally developed by Aronszajn in [Aronszajn, 1950], numerous recent works provide detailed exposition of this topic. Two particular examples are the introductions provided in [Manton et al., 2015] and [Paulsen and Raghupathi, 2016]. We provide a brief outline of the construction of \mathcal{H}_K , and direct the reader to either of [Manton et al., 2015] or [Paulsen and Raghupathi, 2016] for full details. The RKHS \mathcal{H}_K can be constructed as follows.

- Define $\mathcal{H}_0 := \text{Span}(\{K_z \mid z \in \mathcal{M}\})$ where, for $z \in \mathcal{M}$, the function $K_z : \mathcal{M} \rightarrow \mathbb{R}$ is defined by $K_z(v) := K(z, v)$ for $v \in \mathcal{M}$.
 - For $v, z \in \mathcal{M}$ define $\langle K_z, K_v \rangle_{\mathcal{H}_0} := K(z, v)$ and subsequently bilinearly extend $\langle \cdot, \cdot \rangle_{\mathcal{H}_0}$ to the entirety of \mathcal{H}_0 .
 - Prove that $(\mathcal{H}_0, \langle \cdot, \cdot \rangle_{\mathcal{H}_0})$ is an inner product space and that if $f \in \mathcal{H}_0$ and $v \in \mathcal{M}$ then $f(v) = \langle f, K_v \rangle_{\mathcal{H}_0}$.
 - Let \mathcal{H}_K be the collection of functions $f : \mathcal{M} \rightarrow \mathbb{R}$ for which there exists a \mathcal{H}_0 -Cauchy sequence $\{f_n\}_{n=1}^\infty \subset \mathcal{H}_0$ such that $f_n \rightarrow f$ pointwise as $n \rightarrow \infty$.
 - Show that $\langle f, g \rangle_{\mathcal{H}_K} := \lim_{n \rightarrow \infty} \langle f_n, g_n \rangle_{\mathcal{H}_0}$ for \mathcal{H}_0 -Cauchy sequences $\{f_n\}_{n=1}^\infty, \{g_n\}_{n=1}^\infty \subset \mathcal{H}_0$ converging pointwise to f and g respectively is a well-defined inner product on \mathcal{H}_K .
 - Prove that \mathcal{H}_K is complete with respect to this inner product, i.e. that $(\mathcal{H}_K, \langle \cdot, \cdot \rangle_{\mathcal{H}_K})$ is a Hilbert space.
- Moreover, it follows that if $f \in \mathcal{H}_K$ and $v \in \mathcal{M}$ then $f(v) = \langle f, K_v \rangle_{\mathcal{H}_K}$.

The property that for $f \in \mathcal{H}_K$ and $v \in \mathcal{M}$ we have $f(v) = \langle f, K_v \rangle_{\mathcal{H}_K}$ is sometimes called the *reproducing property*. It is worth mentioning that this outlined strategy works under weaker assumptions for the kernel $K : \mathcal{M} \times \mathcal{M} \rightarrow \mathbb{R}$. Indeed it is only required that the kernel be symmetric in the sense that for every $v, u \in \mathcal{M}$ we have $K(u, v) = K(v, u)$, and positive definite in the sense that for every positive integer $n \in \mathbb{Z}_{\geq 1}$ and any $v_1, \dots, v_n \in \mathcal{M}$ and $a_1, \dots, a_n \in \mathbb{R}$ that $\sum_{i=1}^n \sum_{j=1}^n a_i a_j K(v_i, v_j) \geq 0$. Both these properties are evidently satisfied when K is given by a real inner product on a real vector space.

Returning to the setting in which we have embeddings as outlined in (2.11) and the resulting function K defined in (2.12), we observe that we now have a concrete approach to the learning problem. Recall we have a labelled dataset $\Omega = \{(\mathbf{x}_i, y_i) \mid i \in \{1, \dots, M\}\} \subset \mathcal{S}(V) \times \mathbb{R}$. We could use the functions $K_{\mathbf{x}_i} \in \mathcal{H}_K$ for $i \in \{1, \dots, M\}$ as the feature functions. That is, first solve the system

$$y_i = \sum_{k=1}^M a_k K_{\mathbf{x}_k}(\mathbf{x}_i) \quad (2.13)$$

for each $i \in \{1, \dots, M\}$, and then express the observed function f as

$$f(\cdot) := \sum_{k=1}^M a_k K_{\mathbf{x}_k}(\cdot). \quad (2.14)$$

When W is a Hilbert space the so-called *Representer Theorem*, established in [Schölkopf et al., 2001], provides theoretical guarantees regarding the solvability of the equations in (2.13) via gradient descent with respect to a suitable cost function, possibly involving a regularisation term, see [Schölkopf et al., 2001] for full details. However, the equations determined in (2.13) clearly do not require W to be a Hilbert space, and the outlined approach still makes sense under the weaker assumption that W is only a Banach space.

Of course in this setting the function K defined in (2.12) fails to be a kernel in the machine learning sense; indeed it need not be either symmetric or positive definite. Consequently one cannot invoke the Representer Theorem established in [Schölkopf et al., 2001] to provide theoretical guarantees regarding the use of cost function based gradient descent to solve the system of equations (2.13). However, it is in fact known that the Representer Theorem remains valid for certain types of Banach spaces; see [Schlegel, 2021]. Hence there are settings in which just having the assumption of a pair of embeddings, one into a Banach space W and the other into its dual W^* , rather than an embedding into a Hilbert space remains sufficient to provide the theoretical guarantees of the Representer theorem.

A major benefit of the kernel approach outlined above is that 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 M^2 numbers $K(x_i, x_j)$, which is independent of the dimension of W . Assuming K is known, computing these M^2 numbers does not require computing the embedded data points $\Phi(\mathbf{x}_i)$ and $\Psi(\mathbf{x}_i)$ themselves. Consequently there is no restriction on the embeddings that may be considered, with embeddings into infinite dimensional Hilbert 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 $\Omega_{S(V)}$, which becomes problematic for large datasets $\Omega_{S(V)}$.

Unlike the motivating example of using monomials as the feature functions (cf. Subsection 2.1), there is no natural notion of ordering on the set $\{K_{\mathbf{x}_i} \mid i \in \mathcal{I}\}$. For the monomials on $[0, 1]$, namely the set $\{\phi_k : [0, 1] \rightarrow \mathbb{R}, \phi_k(x) := x^k \mid k \in \mathbb{Z}_{\geq 0}\}$, then there is the grading given by the order of the monomials. There is an intrinsic idea that the higher the order the more costly it is to use, and thus it is somehow preferable to use lower order monomials rather than ones of higher order; i.e. using ϕ_2 in place of ϕ_{10} would be less costly. A priori there is no such intrinsic grading on the set of functions $\{K_{\mathbf{x}_i} \mid i \in \{1, \dots, M\}\}$. Given distinct $i, j \in \{1, \dots, M\}$, it is not clear whether there is any benefit between using $K_{\mathbf{x}_i}$ or $K_{\mathbf{x}_j}$ to approximate a function.

2.5. Streams of Increments as Paths

A stream of increments determines a system whose state evolves over time. It may be viewed as a map from an ordered state space to the Banach space V , and concatenating the increments provides an intrinsic way of embedding the stream within a continuous path $\gamma : I \rightarrow V$ for some compact interval $I \subset \mathbb{R}$. Having obtained our streams of increments via the transformation detailed in Subsections 2.2 and 2.3, the continuous path resulting from the concatenation of the entries will be a closed loop based at the origin $0 \in V$. This interpolation of a stream of increments is, in some sense, less *ad-hoc* than the variety of interpolation schemes that can be applied to a stream of values. Examples of several explicit choices of such interpolation schemes, including linear interpolation, rectilinear interpolation, and cubic spline interpolation, can be found in [Chevyrev and Kormilitzin, 2016].

The concatenation of the increments contained within a stream of increments determines an *unparameterised* path in V ; sometimes also referred to as a *shape* or a *contour*. Each entry tells us how the system evolves; but the time over which this evolution happens is determined via a choice of how to view contour. This choice should not influence the affect of the system. Our choice of how to view the contour should not affect the systems response to the contour. We will summarise the effects of a stream of increments by using the signature to succinctly capture the effects of the contour resulting from the concatenation of the entries in the stream.

It is worth remarking that streams of increments can be *augmented* in the same way that streams of values often are. Typically, supposing that $\mathbf{x} = (x_1, \dots, x_k) \in V^k$ for some $k \in \mathbb{Z}_{\geq 1}$, one considers a fixed map $\Theta : V \rightarrow W$ for some Banach space W , and subsequently uses Θ to transform \mathbf{x} to the augmented stream $\Theta(\mathbf{x}) := (\Theta(x_1), \dots, \Theta(x_k)) \in W^k$. The underlying idea is that the augmentation makes it easier to summarise the resulting stream, possibly by explicitly introducing a characteristic that is desirable for the particular model. Evidently the choice of augmentation to use is a modelling problem; it will depend upon the specific task one is aiming to solve.

Once the dataset $\Omega_{\mathcal{S}(V)}$ is augmented by Θ to a dataset $\Omega_{\mathcal{S}(W)} := \{\Theta(\mathbf{x}) \mid \mathbf{x} \in \Omega_{\mathcal{S}(V)}\} \subset \mathcal{S}(W)$, one can consider the regression problem of transitioning from the pairs $\{(\Theta(\mathbf{x}_i), y_i) \mid i \in \{1, \dots, M\}\}$ to a continuous function $h : \Theta(\mathcal{M}) \rightarrow \mathbb{R}$. Then the original desired continuous function $f : \mathcal{M} \rightarrow \mathbb{R}$ would be given by $f := h \circ \Theta$.

We end this section by explicitly introducing some of the more commonly used augmentations as maps $\mathcal{S}(V) \rightarrow \mathcal{S}(W)$ for a choice of Banach space W that is specific to each augmentation.

The *Lead-Lag* augmentation is the map $T_{\text{Lead-Lag}} : \mathcal{S}(V) \rightarrow \mathcal{S}(V \times V)$ defined by

$$T_{\text{Lead-Lag}}(v_1, \dots, v_k) := ((v_1, v_1), (v_2, v_1), (v_2, v_2), \dots, (v_j, v_j), (v_{j+1}, v_j), \dots, (v_k, v_k)). \quad (2.15)$$

It takes a stream with k entries to a stream with $2k$ entries, 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 number of steps between the future being updated and the past being updated can be varied. 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.

For the purpose of an explicit illustration of the Lead-Lag augmentation, take $V := \mathbb{R}$ and consider the stream $\Omega_L := \{1, 6, 3, 9, 5\}$. The basic Lead-Lag augmentation defined in (2.15) 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

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

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

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

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

$$\left\{ \begin{pmatrix} v_1 \\ 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} v_2 \\ v_1 \\ 0 \\ \vdots \\ 0 \\ 0 \end{pmatrix}, \dots, \begin{pmatrix} v_{d+1} \\ v_d \\ v_{d-1} \\ \vdots \\ v_2 \\ v_1 \end{pmatrix}, \begin{pmatrix} v_{d+2} \\ v_{d+1} \\ v_d \\ \vdots \\ v_3 \\ v_2 \end{pmatrix}, \dots, \begin{pmatrix} v_k \\ v_{k-1} \\ v_{k-2} \\ \vdots \\ v_{k-d+1} \\ v_{k-d} \end{pmatrix}, \begin{pmatrix} 0 \\ v_k \\ v_{k-1} \\ \vdots \\ v_{k-d+2} \\ v_{k-d+1} \end{pmatrix}, \dots, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ v_k \\ v_{k-1} \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \\ \vdots \\ 0 \\ v_k \end{pmatrix} \right\} \quad (2.16)$$

in $\mathcal{S}(V^{d+1})$. Returning to the dataset Ω_L , applying the no future pause augmentation defined in (2.16) with $d = 2$ results in the new stream

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

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.

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

$$T_{\text{Time}}((v_1, \dots, v_k)) := ((v_1, t_0), \dots, (v_j, t_j), \dots, (v_k, t_k)) \quad (2.17)$$

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

$$T_{\text{Time-diff}}((v_1, \dots, v_k)) := ((v_1, t_0), \dots, (v_j, t_j - t_{j-1}), \dots, (v_k, t_k - t_{k-1})) \quad (2.18)$$

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 consisting of k entries to streams once again consisting of k entries.

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

$$T_{\text{Time}}(\Omega_T) = \{(1, 0), (5, 1), (2, 3), (9, 6), (7, 8), (6, 12)\} \in \mathcal{S}(\mathbb{R}^2) \quad (2.19)$$

and

$$T_{\text{Time-diff}}(\Omega_T) = \{(1, 0), (5, 1), (2, 2), (9, 3), (7, 2), (6, 4)\} \in \mathcal{S}(\mathbb{R}^2). \quad (2.20)$$

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

$$T_{\text{inv}}(v_1, \dots, v_k) := ((v_1, 1), \dots, (v_j, 1), \dots, (v_k, 1), (v_k, 0), (0, 0)) \quad (2.21)$$

and taking streams with k entries to streams with $k + 2$ entries. We illustrate this augmentation by applying it to both 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

$$T_{\text{inv}}(\Omega_{ir1}) = \{(1, 1), (3, 1), (4, 1), (8, 1), (9, 1), (9, 0), (0, 0)\} \in \mathcal{S}(\mathbb{R}^2) \quad (2.22)$$

and

$$T_{\text{inv}}(\Omega_{ir2}) = \left\{ \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 3 \\ 4 \\ 1 \end{pmatrix}, \begin{pmatrix} 4 \\ 6 \\ 1 \end{pmatrix}, \begin{pmatrix} 5 \\ 9 \\ 1 \end{pmatrix}, \begin{pmatrix} 7 \\ 10 \\ 1 \end{pmatrix}, \begin{pmatrix} 7 \\ 10 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} \right\} \in \mathcal{S}(\mathbb{R}^3). \quad (2.23)$$

This augmentation is utilised in the transformation for converting a stream of values to a stream of increments that is detailed in Subsection 2.3. A particular consequence of this is that the signature of the resulting stream of increments will capture information related to the norm/size of the entries in the original stream of values (cf. Section 3).

2.6. Signature Involvement

The finite dataset $\Omega_{\mathcal{S}(V)}$ is a subset of the set of streams on V denoted by $\mathcal{S}(V)$. Thus given a stream $\mathbf{x} \in \mathcal{S}(V)$, we seek to summarise \mathbf{x} (throw away irrelevant information) so as to capture its effect. Recalling our assumption that the streams we consider are formed of incremental data, we observe that each stream can be associated with the contour/shape determined by the concatenation of the increments. Given a stream of increments $\mathbf{x} \in \mathcal{S}(V)$ we denote the contour arising as the concatenation of the increments as $\Gamma_{\mathbf{x}}$. The approach followed in this article is to summarise a stream of increments $\mathbf{x} \in \mathcal{S}(V)$ by summarising the contour $\Gamma_{\mathbf{x}}$.

It is instructive to realise that the classical mathematical description of paths is ill-suited to providing

such a summary. Assume that we have a parameterisation $\gamma : I \rightarrow V$ of the contour $\Gamma_{\mathbf{x}}$ where $I \subset \mathbb{R}$ is a compact interval. Then the classical description of the path γ fails to adequately compress the information it contains. 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. For example, it could be critical from a diagnostic perspective to know whether a patients heart rate or their breathing rate increased 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 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 want to summarise the contour $\Gamma_{\mathbf{x}}$ associated to the stream of increments \mathbf{x} . The contour $\Gamma_{\mathbf{x}}$ is independent of the choice of parameterisation we make to view it. This reflects the fact that the response of a system to a given stream \mathbf{x} is not affected by the speed at which the provided sequential increments happen. The features of the stream induce the same response regardless of how quickly we choose to experience them. For example, consider writing the character '3'. The effect of this stream 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 view the stream given by 'writing a figure 3' does *not* change its effect in producing a '3' on our page.

Consequently, if we want to summarise $\Gamma_{\mathbf{x}}$ by summarising a parameterisation $\gamma : I \rightarrow V$ of $\Gamma_{\mathbf{x}}$ defined on a compact interval $I \subset \mathbb{R}$, we must summarise γ in a manner that is invariant under reparameterisation. That is, we must quotient out the symmetry induced by the group of reparameterisations so that the resulting summary of the contour $\Gamma_{\mathbf{x}}$ depends only on the contour $\Gamma_{\mathbf{x}}$ and *not* on any particular choices made for the purposes of viewing it. The information extracted from the contour $\Gamma_{\mathbf{x}}$ should not depend on how we choose to access $\Gamma_{\mathbf{x}}$.

The freedom of reparameterisation makes embedding a stream into a continuous path a reasonable idea. Each increment in a stream \mathbf{x} represents a change in the state of the system. Initially these changes appear to happen instantaneously, seemingly giving rise to a discontinuous path. Hence it appears that the contour $\Gamma_{\mathbf{x}}$ arising from the concatenation of these changes in a continuous manner is not a good representation of the stream \mathbf{x} .

However, under the assumption that reparameterisation does not affect the underlying system, 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 associating each stream $\mathbf{x}_i \in \Omega_{\mathcal{S}(V)}$ with its corresponding contour $\Gamma_{\mathbf{x}_i}$, our problem has been transformed to seeking to learn a function that is an effect of these unparameterised 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 γ 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

$$\exp(x) = \sum_{k=0}^{\infty} \frac{x^k}{k!}. \quad (2.24)$$

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.

2.7. Tensor Algebra and Signature

Let V be a d -dimensional Banach space. The spaces of *formal polynomials* and *formal power series* over V are given by

$$T(V) := \bigoplus_{k=0}^{\infty} V^{\otimes k} \quad \text{and} \quad T((V)) := \prod_{k=0}^{\infty} V^{\otimes k} \quad (2.25)$$

respectively where $V^{\otimes 0} := \mathbb{R}$. Both addition, scalar multiplication, and the tensor product \otimes extend to $T((V))$ as follows. Suppose that $A = \prod_{k=0}^{\infty} a_k \in T((V))$, $B = \prod_{k=0}^{\infty} b_k \in T((V))$, and $\lambda \in \mathbb{R}$. Then we define the operations

$$A + B := \prod_{k=0}^{\infty} a_k + b_k \quad \text{and} \quad A \otimes B := \prod_{n=0}^{\infty} \sum_{k=0}^n a_k \otimes b_{n-k} \quad (2.26)$$

and recall the natural action of \mathbb{R} on $T((V))$ given by

$$\lambda A := \prod_{k=0}^{\infty} \lambda a_k. \quad (2.27)$$

The operations defined in (2.26) and the action of \mathbb{R} specified in (2.27) also make sense for elements in $T(V)$.

When equipped with the operations in (2.26) and the action of \mathbb{R} defined in (2.27) $T((V))$ becomes a real, non-commutative unital algebra, with unit $\mathbf{1} = (1, 0, 0, \dots)$, called the *Tensor algebra* of V . It is within the tensor algebra that the signature of a path $x : [a, b] \rightarrow V$ lives. Before detailing this, we first consider equipping $T((V))$ with a norm.

For convenience, for each $n \in \mathbb{Z}_{\geq 0}$, let $\pi_n : T((V)) \rightarrow V^{\otimes n}$ denote the projection map. Moreover, for each $n \in \mathbb{Z}_{\geq 0}$, we may consider the *truncated tensor algebra*

$$T^{(n)}(V) := \prod_{k=0}^n V^{\otimes k} \quad (2.28)$$

and denote the projection $T((V)) \rightarrow T^{(n)}(V)$ by Π_n . One particular way to equip the tensor algebra $T((V))$ with a norm would be to choose *admissible tensor norms*, in the sense originating in [Schatten, 1985], for all the tensor powers of V . For clarity, we give a precise definition below.

Definition 2.1 (Admissible Tensor Norms). Let V be a Banach space. The tensor powers of V are equipped with admissible norms if for each integer $n \in \mathbb{Z}_{\geq 1}$ we have chosen a norm on $V^{\otimes n}$ such that the following conditions are satisfied.

- (A) For each $n \in \mathbb{Z}_{\geq 1}$ the norm $\|\cdot\|_{V^{\otimes n}}$ on $V^{\otimes n}$ is invariant under the action of the symmetric group S_n on $V^{\otimes n}$. To elaborate, given $\rho \in S_n$ and $a_1 \otimes \dots \otimes a_n \in V^{\otimes n}$, then $\rho(a_1 \otimes \dots \otimes a_n) := a_{\rho(1)} \otimes \dots \otimes a_{\rho(n)}$, and the action is extended to the entirety of $V^{\otimes n}$ by linearity. Then it is required that for every $v \in V^{\otimes n}$ and every $\rho \in S_n$ we have $\|\rho(v)\|_{V^{\otimes n}} = \|v\|_{V^{\otimes n}}$.
- (B) For any $n, m \in \mathbb{Z}_{\geq 1}$ and any $v \in V^{\otimes n}$ and $w \in V^{\otimes m}$ we have $\|v \otimes w\|_{V^{\otimes(n+m)}} \leq \|v\|_{V^{\otimes n}} \|w\|_{V^{\otimes m}}$.

- (C) For any integers $n, m \in \mathbb{Z}_{\geq 1}$ and for any dual elements $\phi \in (V^{\otimes n})^*$ and $\sigma \in (V^{\otimes m})^*$ we have $\|\phi \otimes \sigma\|_{(V^{\otimes(n+m)})^*} \leq \|\phi\|_{(V^{\otimes n})^*} \|\sigma\|_{(V^{\otimes m})^*}$. Here, given any $k \in \mathbb{Z}_{\geq 1}$, the norm $\|\cdot\|_{(V^{\otimes k})^*}$ denotes the dual-norm induced by $\|\cdot\|_{V^{\otimes k}}$.

It turns out (see [Ryan, 2002]) that having *both* the inequalities stated in Definition 2.1 (B) and (C) respectively ensures that we in fact have equality in both estimates. Hence if the tensor powers of V are equipped with admissible norms, we have equality in both the inequalities stated in Definition 2.1 (B) and (C) respectively.

Two common choices are the *projective* and *injective* tensor norms. The projective tensor norm is defined, for $n \geq 2$, on $V^{\otimes n}$ by setting, for $v \in V^{\otimes n}$,

$$\|v\|_{\text{proj},n} := \inf \left\{ \sum_{i=1}^{\infty} \prod_{j=1}^n \|a_{j_i}\|_V \mid v = \sum_{i=1}^{\infty} a_{1_i} \otimes \dots \otimes a_{n_i} \text{ and } \sum_{i=1}^{\infty} \prod_{j=1}^n \|a_{j_i}\|_V < \infty \right\}. \quad (2.29)$$

The injective tensor norm is defined, for $n \geq 2$, on $V^{\otimes n}$ by setting, for $v \in V^{\otimes n}$,

$$\|v\|_{\text{inj},n} := \sup \left\{ |\varphi_1 \otimes \dots \otimes \varphi_n(v)| \mid \varphi_1, \dots, \varphi_n \in V^* \text{ and } \|\varphi_1\|_{V^*} = \dots = \|\varphi_n\|_{V^*} = 1 \right\}. \quad (2.30)$$

As observed in [Schatten, 1985], the injective and projective tensor norms are the prototype examples of admissible tensor norms in the following sense. Given any arbitrary choice of admissible tensor norms on the tensor powers of V in the sense of Definition 2.1, then for every $n \in \mathbb{Z}_{\geq 2}$, if $\|\cdot\|_{V^{\otimes n}}$ denotes the norm chosen for $V^{\otimes n}$, we have for every $v \in V^{\otimes n}$ that (cf. Proposition 2.1 in [Ryan, 2002])

$$\|v\|_{\text{inj},n} \leq \|v\|_{V^{\otimes n}} \leq \|v\|_{\text{proj},n}. \quad (2.31)$$

Any choice of admissible tensor norms for the tensor powers of V could be used to define a norm on the tensor algebra $T((V))$. For example, we could define

$$\|v\|_{T((V))} := \sum_{n=0}^{\infty} \|\pi_n(v)\|_{V^{\otimes n}} \quad (2.32)$$

for $v \in T((V))$ with the choice of the norm on $V^{\otimes 0} = \mathbb{R}$ as the usual absolute value.

For separable real Banach spaces V it is known that one may define an inner product on V so that the Hilbert space resulting from the completion of V with respect to this inner product is closely related to the original Banach space structure on V ; see [Kuelbs, 1970]. In the case that V is finite dimensional, as we are assuming, this can be achieved in the following way. We begin by letting $\mathcal{B} = \{v_1, \dots, v_d\}$ be a basis for V , and subsequently denote the induced dual basis of V^* by $\mathcal{B}^* = \{v_1^*, \dots, v_d^*\}$. For each $n \in \mathbb{Z}_{\geq 2}$ the basis \mathcal{B} determines the basis

$$\mathcal{B}^{\otimes n} := \left\{ v_{\mathbf{K}} = v_{k_1} \otimes \dots \otimes v_{k_n} \mid \mathbf{K} = (k_1, \dots, k_n) \in \{1, \dots, d\}^n \right\} \quad (2.33)$$

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

$$(\mathcal{B}^*)^{\otimes n} := \left\{ v_{\mathbf{K}}^* = v_{k_1}^* \otimes \dots \otimes v_{k_n}^* \mid \mathbf{K} = (k_1, \dots, k_n) \in \{1, \dots, d\}^n \right\} \quad (2.34)$$

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

We first use the basis \mathcal{B} to define an inner product on V . To do so we set

$$\langle v_i, v_j \rangle_V := \delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \quad (2.35)$$

and extending bilinearly to the whole of V . Given any $n \in \mathbb{Z}_{\geq 2}$ we use the basis $\mathcal{B}^{\otimes n}$ to define an inner

product on $V^{\otimes n}$. We set

$$\langle v_{i_1} \otimes \dots \otimes v_{i_n}, v_{j_1} \otimes \dots \otimes v_{j_n} \rangle_{V^{\otimes n}} := \prod_{k=1}^n \langle v_{i_k}, v_{j_k} \rangle_V = \delta_{i_1 j_1} \dots \delta_{i_n j_n} \quad (2.36)$$

and again extending bilinearly to the whole of $V^{\otimes n}$. Finally, we define an inner product on the tensor algebra $T((V))$ by setting

$$\langle A, B \rangle_{T((V))} = \sum_{n=0}^{\infty} \langle \pi_n(A), \pi_n(B) \rangle_{V^{\otimes n}} \quad (2.37)$$

for $A, B \in T((V))$ where we choose the inner product on $V^{\otimes 0} = \mathbb{R}$ to be given by usual multiplication. Let $\mathcal{H}(V) := \overline{T((V))}$ denote the Hilbert space given by the completion of the tensor algebra $T((V))$ with respect to this inner product. Throughout this article we assume that the tensor algebra $T((V))$ has been equipped with an inner product in this manner.

The choices of inner products on the tensor powers of V interact with the induced dual bases in the expected way. To be more precise, first observe that for every $n \in \mathbb{Z}_{\geq 1}$ we have that $(V^*)^{\otimes n} = (V^{\otimes n})^*$. Consequently, given $n \in \mathbb{Z}_{\geq 1}$, if $\mathbf{K} \in \{1, \dots, d\}^n$ then the element $v_{\mathbf{K}}^* \in \mathcal{B}^* \subset (V^{\otimes n})^*$ is given by $v_{\mathbf{K}}^*(\cdot) = \langle \cdot, v_{\mathbf{K}} \rangle_{V^{\otimes n}}$.

Momentarily ignoring regularity issues, the *signature* of a continuous path $x : [a, b] \rightarrow V$ is the solution to the universal differential equation

$$dS_{a,t}(x) = S_{a,t}(x) \otimes dx_t \quad \text{with} \quad S_{a,a}(x) = \mathbf{1} = (1, 0, 0, \dots) \in T((V)), \quad (2.38)$$

formalising that the signature $S(x)$ is the response of the exponential nonlinear system to the input x . The subsequent Subsections 2.8 and 2.9 are dedicated to examining the regularity issues related to the differential equation defined in (2.38).

We end this subsection by defining certain subsets of the tensor algebra $T((V))$ and particular operations on elements within the tensor algebra $T((V))$ that will be useful in the later subsections. First, we introduce the notation that $\tilde{T}((V))$ will denote the collection of elements $\mathbf{a} \in T((V))$ for which $\pi_0(\mathbf{a}) = 1 \in \mathbb{R}$, and $T_{>0}((V))$ will denote the collection of elements $\mathbf{a} \in T((V))$ for which $\pi_0(\mathbf{a}) > 0$.

By using the power series of the real-valued functions $x \mapsto e^x$ and $x \mapsto \log(1+x)$, we define the *exponential* \exp and *logarithm* \log of elements within the tensor algebra $T((V))$.

Definition 2.2 (Exponential and Logarithm; Definition 2.20 in [Lyons et al., 2007]). The *exponential* map on the tensor algebra $T((V))$ is the map $\exp : T((V)) \rightarrow \tilde{T}((V))$ is defined for $\mathbf{a} \in T((V))$ by

$$\exp(\mathbf{a}) := \sum_{n=0}^{\infty} \frac{\mathbf{a}^{\otimes n}}{n!}. \quad (2.39)$$

The *logarithm* map $\log : T_{>0}((V)) \rightarrow T((V))$ is defined for $\mathbf{a} \in T_{>0}((V))$ by

$$\log(\mathbf{a}) := \log(\pi_0(\mathbf{a})) + \sum_{n=1}^{\infty} \frac{(-1)^n}{n} \left(\mathbf{1} - \frac{\mathbf{a}}{\pi_0(\mathbf{a})} \right)^{\otimes n} \quad (2.40)$$

where we recall that π_0 is the projection map $T((V)) \rightarrow \mathbb{R} =: V^{\otimes 0}$.

The series defining the map \exp in 2.39 is convergent (see Lemma 2.19 in [Lyons et al., 2007]). Further, the series defining the map \log in 2.40 is purely algebraic involving only finitely many terms at each level. Consequently, it does not depend on the norms on the tensor powers of V . Moreover, if we let B_1 denote the collection of elements $\mathbf{a} \in T((V))$ with $\pi_0(\mathbf{a}) = 0$, then $\exp : B_1 \rightarrow \tilde{T}((V))$ and $\log : \tilde{T}((V)) \rightarrow B_1$ are both one-to-one and the inverse of one and other; see Lemma 2.21 in [Lyons et al., 2007].

The tensor algebra $T((V))$ carries a Lie bracket $[\cdot, \cdot] : T((V)) \times T((V)) \rightarrow T((V))$ defined for $\mathbf{a}, \mathbf{e} \in T((V))$ by

$$[\mathbf{a}, \mathbf{e}] := \mathbf{a} \otimes \mathbf{e} - \mathbf{e} \otimes \mathbf{a}. \quad (2.41)$$

Given linear subspaces $F_1, F_2 \subset T((V))$ we define

$$[F_1, F_2] := \text{Span} \left(\{ [\mathbf{a}, \mathbf{e}] \mid \mathbf{a} \in F_1 \text{ and } \mathbf{e} \in F_2 \} \right). \quad (2.42)$$

Then we may define a sequence $\{L_n\}_{n=0}^\infty$ of linear subspaces of $T((V))$ recursively by

$$L_0 := \{0\} \subset \mathbb{R} =: V^{\otimes 0}, \quad L_1 := V \text{ and, for every } n \in \mathbb{Z}_{\geq 1} \quad L_{n+1} := [V, L_n] \subset V^{\otimes(n+1)}. \quad (2.43)$$

For each $n \in \mathbb{Z}_{\geq 0}$ the subspace $L_n \subset V^{\otimes n}$ defined in (2.43) is called the space of *homogeneous Lie polynomials* of degree n . The sequence $\{L_n\}_{n=0}^\infty$ enables us to define the space of *Lie formal series* over V .

Definition 2.3 (Lie formal series over V ; Definition 2.22 in [Lyons et al., 2007]). The space of *Lie formal series* over V is the subspace $\mathcal{L}((V)) \subset T((V))$ given by

$$\mathcal{L}((V)) := \left\{ \mathbf{1} \in T((V)) \mid \text{for every } n \in \mathbb{Z}_{\geq 0} \text{ we have } \pi_n(\mathbf{1}) \in L_n \right\}. \quad (2.44)$$

Moreover, for each $n \in \mathbb{Z}_{\geq 0}$ we define the *Lie polynomials* of degree n to be the subset $\mathcal{L}^{(n)}(V) := \Pi_n(\mathcal{L}((V)))$.

Finally we define the set of *group-like* elements in $\tilde{T}((V))$.

Definition 2.4 (Group-like elements). The *group-like* elements are those in the subset $G^{(*)} \subset \tilde{T}((V))$ defined by

$$G^{(*)} := \exp \left(\mathcal{L}((V)) \right). \quad (2.45)$$

Further, for each integer $n \in \mathbb{Z}_{\geq 0}$, the *free nilpotent group of order n* is defined as $G^{(n)} := \Pi_n(G^{(*)})$.

A particular consequence of (2.45) is that $\mathcal{L}((V)) = \log(G^{(*)})$. Moreover, as the name suggests, the group-like elements $G^{(*)}$ form a group under the tensor product operation \otimes . We will later see that the signature of a large class of paths lie within this group (cf. Subsections 2.8 and 2.9).

2.8. Signature of Paths with Finite p-Variation

In this subsection we cover the signature of a continuous path $x : [a, b] \rightarrow V$ having finite p -variation for some $p \in [1, 2)$. We begin by introducing the notion of p -variation regularity.

If $p \in \mathbb{R}_{\geq 1}$ then the p -variation of a continuous path $x : [a, b] \rightarrow V$ is defined to be

$$\|x\|_{p,[a,b]} := \left(\sup_{a \leq t_0 < t_1 < \dots < t_r \leq b} \left\{ \sum_{j=0}^{r-1} \|x_{t_{j+1}} - x_{t_j}\|_V^p \right\} \right)^{\frac{1}{p}}. \quad (2.46)$$

We denote the collection of continuous paths $[a, b] \rightarrow V$ with finite p -variation by $\mathcal{V}^p([a, b], V)$. We can equip $\mathcal{V}^p([a, b], V)$ with a norm called the p -variation norm by defining, for $x \in \mathcal{V}^p([a, b], V)$,

$$\|x\|_{\mathcal{V}^p([a,b],V)} := \|x\|_{p,[a,b]} + \sup_{t \in [a,b]} \{ \|x_t\|_V \}. \quad (2.47)$$

If $p, q \in \mathbb{R}_{\geq 1}$ with $p \leq q$ then $\mathcal{V}^1([a, b], V) \subset \mathcal{V}^p([a, b], V) \subset \mathcal{V}^q([a, b], V) \subset C^0([a, b], V)$ (see Proposition 1.7 in [Lyons et al., 2007], for example). Further, if $\alpha \in (0, 1]$ and $y \in C^0([a, b], V)$ is α -Hölder continuous in the sense that there exists a constant $C > 0$ such that $\|y_t - y_s\|_V \leq C|t - s|^\alpha$ whenever $s, t \in [a, b]$, then y has finite $1/\alpha$ -variation, i.e. $y \in \mathcal{V}^{1/\alpha}([a, b], V)$. Consequently, a typical Brownian path $B : [a, b] \rightarrow \mathbb{R}$ has finite p -variation for every $p > 2$, i.e. $B \in \mathcal{V}^p([a, b], \mathbb{R})$. It is however known that the 2-variation of B is almost surely infinite [Freedman, 2012]. However, if we instead view the Brownian path as a path into L^2 rather than \mathbb{R} , then its 2-variation on $[a, b]$ is finite; see [Lyons et al., 2007], for example.

If a continuous path $x : [a, b] \rightarrow V$ is in $\mathcal{V}^p([a, b], V)$ for a real number $1 \leq p < 2$ then its signature (i.e. the solution to the differential equation (2.38)) can be constructed explicitly using the *Young integral*, which is a continuous extension of the Stieltjes integral. Given Banach spaces V and W , the Young integral determines a way for a path $[a, b] \rightarrow \mathbf{L}(V, W)$ to be integrated along a path $[a, b] \rightarrow V$. The notation $\mathbf{L}(V, W)$ denotes the space of linear operators $V \rightarrow W$. To be more precise, let $X : [a, b] \rightarrow V$ and $Y : [a, b] \rightarrow \mathbf{L}(V, W)$. Given a partition $\mathcal{D} = (a = t_0, t_1, \dots, t_r = b)$ of $[a, b]$ we let $|\mathcal{D}| := \min \{t_{i+1} - t_i \mid i \in \{0, \dots, r-1\}\}$. We first define

$$\int_{\mathcal{D}} Y dX := \sum_{i=0}^{r-1} Y_{t_i} [X_{t_{i+1}} - X_{t_i}]. \quad (2.48)$$

Then, given $t \in [a, b]$, the Young integral of Y along X over the interval $[a, t]$ is defined as

$$\int_a^t Y_s dX_s := \lim_{j \rightarrow \infty} \int_{\mathcal{D}_j} Y dX \quad (2.49)$$

where $\{\mathcal{D}_j\}_{j=1}^{\infty}$ is a sequence of partitions of $[a, t]$ with $|\mathcal{D}_j| \rightarrow 0$ as $j \rightarrow \infty$. When $p, q \in \mathbb{R}_{\geq 1}$ with $\frac{1}{p} + \frac{1}{q} > 1$, $X \in \mathcal{V}^p([a, b], V)$, and $Y \in \mathcal{V}^q([a, b], \mathbf{L}(V, W))$ it can be established that the mapping $t \mapsto \int_a^t Y_s dX_s$ defined in (2.49) is a well-defined path $[a, b] \rightarrow W$ with finite p -variation; see Theorem 1.16 in [Lyons et al., 2007].

Using the Young integral, if $p \in [1, 2)$ and $x \in \mathcal{V}^p([a, b], V)$, then the signature of x is given by the iterated integrals of x . That is, $S_{a,t}(x) = \prod_{n=0}^{\infty} S_{a,t}^n(x)$ where

$$S_{a,t}^0(x) \equiv 1 \quad \text{and, for } n \in \mathbb{Z}_{\geq 1}, \quad S_{a,t}^n(x) = \int_{a < t_1 < \dots < t_n < t} dx_{t_1} \otimes \dots \otimes dx_{t_n}. \quad (2.50)$$

The path $t \mapsto S_{a,t}(x)$ is the unique solution to the differential equation stated in (2.38) for the path x ; see Lemma 2.10 in [Lyons et al., 2007], for example.

There are estimates known for the norms of the components of the signature $S_{a,t}(x)$, i.e. there are known upper bounds for the quantities $\|S_{a,t}^n(x)\|_{V^{\otimes n}}$. In order to state these estimates we introduce the notion of a *control* on the interval $[a, b]$. First, we define $\Delta_{[a,b]} := \{(s, t) \mid a \leq s \leq t \leq b\}$. A control on $[a, b]$ is a continuous function $\omega : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ that is super-additive, in the sense that for every $s, t, u \in [a, b]$ with $s \leq u \leq t$ we have $\omega(s, u) + \omega(u, t) \leq \omega(s, t)$, and vanishes on the diagonal, in the sense that for every $t \in [a, b]$ we have $\omega(t, t) = 0$.

Raising the p -variation of a path x to the power of p provides an example of a control. That is, the function $\omega_x : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ defined for $(s, t) \in \Delta_{[a,b]}$ by $\omega_x(s, t) := \|x\|_{p, [s, t]}^p$ is a control. Given a general control $\omega : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$, the p -variation of a path $x : [a, b] \rightarrow V$ is controlled by ω if for every $(s, t) \in \Delta_{[a,b]}$ we have $\|x\|_{p, [s, t]} \leq \omega(s, t)^{1/p}$.

Now suppose that $p \in [1, 2)$ and that $x \in \mathcal{V}^p([a, b], V)$. Then there exists a control $\omega : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ such that for every $n \in \mathbb{Z}_{\geq 1}$ and every $(s, t) \in \Delta_{[a,b]}$ we have

$$\|S_{s,t}^n(x)\|_{V^{\otimes n}} \leq \frac{\omega(s, t)^n}{\Gamma(1 + \frac{n}{p})} \quad (2.51)$$

where Γ denotes the Gamma function. In particular, for positive real numbers $a > 0$ we have that $\Gamma(a) := \int_0^{\infty} r^{a-1} e^{-r} dr$, and for integers $m \in \mathbb{Z}_{\geq 1}$ we have $\Gamma(1 + m) = m!$. This latter equality motivates the commonly used notation of writing $(n/p)!$ for $\Gamma(1 + n/p)$ in (2.51).

When $p = 1$ one can take $\omega(s, t) := \omega_x(s, t) = \|x\|_{1, [s, t]}$ and establish the estimate (2.51) via a direct calculation; see Proposition 2.2 in [Lyons et al., 2007]. Consequently, in the case that $p = 1$ we have the estimate that for every $n \in \mathbb{Z}_{\geq 1}$ and every $(s, t) \in \Delta_{[a,b]}$ that

$$\|S_{s,t}^n(x)\|_{V^{\otimes n}} \leq \frac{\|x\|_{1, [s, t]}^n}{n!}. \quad (2.52)$$

One way to arrive at the estimates (2.51) for $p \in (1, 2)$ is to invoke the *Extension Theorem* for multiplicative functionals $\Delta_{[a,b]} \rightarrow T((V))$ appearing as Theorem 3.7 in [Lyons et al., 2007]. A definition of what is meant by multiplicative functional $\Delta_{[a,b]} \rightarrow T((V))$ is provided below (cf. (2.53)) and can also be found in Definition 2.6, whilst a variant of the *Extension Theorem* appearing as Theorem 3.7 in [Lyons et al., 2007] is stated as Theorem 2.9 in Subsection 2.9.

This Young integral approach does not work when $p \geq 2$; however, the $p = 2$ threshold is *not* an artefact due to some particular limitation of the Young integral approach. In fact no continuous extension of the Stieltjes integral can be rich enough to handle Brownian paths. Further, the seemingly simple function mapping a path to the area enclosed by the path fails to be continuous in the 2-variation norm. A detailed discussion of these particular issues can be found in Section 1.5 on [Lyons et al., 2007].

In the remainder of this subsection we state some properties of signatures of paths with finite p -variation for $p \in [1, 2)$ that are particularly useful from our machine learning perspective. For this purpose we let $p \in [1, 2)$ and fix a path $x \in \mathcal{V}^p([a, b], V)$.

The signature $S(x)$ is a *multiplicative* functional $\Delta_{[a,b]} \rightarrow T((V))$. By multiplicative, we mean that Chen's identity holds; namely,

$$S_{s,u}(x) \otimes S_{u,t}(x) = S_{s,t}(x) \quad (2.53)$$

whenever $a \leq s \leq u \leq t \leq b$. A consequence of (2.53) 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

$$(x * y)_t := \begin{cases} x_t & \text{if } a \leq t \leq b \\ x_b - y_b + y_t & \text{if } b < t \leq c. \end{cases} \quad (2.54)$$

It then follows from (2.53) that

$$S_{a,c}(x * y) = S_{a,b}(x) \otimes S_{b,c}(y). \quad (2.55)$$

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

$$S_{a,b}^{(n)}(x) := \Pi_n(S_{a,b}(x)). \quad (2.56)$$

It follows that $t \mapsto S_{a,t}^{(n)}(x)$ is the solution to a truncated version of the differential equation stated in (2.38) for the path x . To be precise, the mapping $[a, b] \rightarrow T^{(n)}(V)$ given by $t \mapsto S_{a,t}^{(n)}(x)$ solves the differential equation

$$dS_{a,t}^{(n)}(x) = S_{a,t}^{(n)}(x) \otimes dx_t \quad \text{with} \quad S_{a,a}^{(n)} = \Pi_n(\mathbf{1}) = (1, 0, \dots, 0) \in T^{(n)}(V). \quad (2.57)$$

A detailed proof of this may be found in Chapter 7 of [Friz and Victoir, 2010], for example.

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

$$S_{a,b}(x)^{\mathbf{K}} := \langle S_{a,b}(x), v_{\mathbf{K}} \rangle_{T((V))} = v_{\mathbf{K}}^*(S_{a,b}^{(n)}(x)). \quad (2.58)$$

The name is justified by observing that in the case that the coefficients of the signature are given by (2.50), we see via (2.36) that

$$S_{a,b}(x)^{\mathbf{K}} = \int \cdots \int_{a \leq t_1 \leq \dots \leq t_n \leq b} \langle dx_{t_1}, v_{k_1} \rangle_V \cdots \langle dx_{t_n}, v_{k_n} \rangle_V. \quad (2.59)$$

The product of two coordinate iterated integrals $S_{a,b}(x)^{\mathbf{K}}$ and $S_{a,b}(x)^{\mathbf{L}}$ yields a quadratic form. 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

$$(i_1, \dots, i_n) \sqcup (j_1, \dots, j_p) := \sum_{\sigma \in \text{Shuff}(n,p)} \left(a_{\sigma^{-1}(1)}, \dots, a_{\sigma^{-1}(n+p)} \right) \quad (2.60)$$

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

$$\text{Shuff}(n, p) := \{\sigma \in S_{n+p} \mid \sigma(1) < \dots < \sigma(n) \text{ and } \sigma(n+1) < \dots < \sigma(n+p)\}. \quad (2.61)$$

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 [Lyons et al., 2007]). That is,

$$S_{a,b}(x)^{\mathbf{K}} S_{a,b}(x)^{\mathbf{L}} = S_{a,b}(x)^{\mathbf{K} \sqcup \mathbf{L}}. \quad (2.62)$$

One consequence of (2.62) 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 (2.62) establishes that the range is closed under multiplication. Hence the range of the signature forms a group in the tensor algebra.

In fact, the group formed by the range of the signature is a subgroup of the group given by the *group-like* elements $G^{(*)}$ introduced in Definition 2.4. In order to formalise this, we first extend the shuffle product to a binary operation on $T(V^*)$.

Given elements $\mathbf{e}^*, \mathbf{f}^* \in T(V^*)$ we define their *shuffle product* $\mathbf{e}^* \sqcup \mathbf{f}^*$ as an element in $T(V^*)$ as follows. If $\mathbf{e}^* = v_{\mathbf{K}}^*$ and $\mathbf{f}^* = v_{\mathbf{L}}^*$ for finite words $\mathbf{K} = (i_1, \dots, i_n) \in \{1, \dots, d\}^n$ and $\mathbf{L} = (j_1, \dots, j_p) \in \{1, \dots, d\}^p$, then we define

$$v_{\mathbf{K}}^* \sqcup v_{\mathbf{L}}^* := \sum_{\sigma \in \text{Shuff}(n, p)} v_{\left(a_{\sigma^{-1}(1)}, \dots, a_{\sigma^{-1}(n+p)}\right)}^* \quad (2.63)$$

where $(a_1, \dots, a_{n+p}) := (i_1, \dots, i_n, j_1, \dots, j_p)$. The definition in (2.63) is then extended by linearity to the entirety of $T(V^*)$.

The result of Theorem 2.15 in [Lyons et al., 2007] can be stated as follows. Let $S(\mathcal{V}^p([a, b], V)) \subset \tilde{T}((V))$ denote the collection of elements in the tensor algebra $T((V))$ that arise as the signature of a path $[a, b] \rightarrow V$ with finite p -variation, i.e. the signature of an element in $\mathcal{V}^p([a, b], V)$. Then whenever $\mathbf{e}^*, \mathbf{f}^* \in T(V^*)$ and $\mathbf{a} \in S(\mathcal{V}^p([a, b], V)) \subset \tilde{T}((V))$ we have that

$$\mathbf{e}^* \sqcup \mathbf{f}^*(\mathbf{a}) = \mathbf{e}^*(\mathbf{a}) \mathbf{f}^*(\mathbf{a}) \quad (2.64)$$

It turns out that the subset of $\tilde{T}((V))$ where the identity (2.64) holds for all $\mathbf{e}^*, \mathbf{f}^* \in T(V^*)$ is a group (see Lemma 2.17 in [Lyons et al., 2007]). Moreover, one can then give the following equivalent definition of the group-like elements $G^{(*)}$.

Lemma 2.5 (Equivalent definition of group-like; Variant of Theorem 2.23 in [Lyons et al., 2007]). *Consider an element $\mathbf{a} \in \tilde{T}((V))$ and consider $T(V^*)$, the space of formal polynomials over V^* , to be equipped with the shuffle product. Define the evaluation map $ev_{\mathbf{a}} : T(V^*) \rightarrow \mathbb{R}$ by $ev_{\mathbf{a}}(\mathbf{e}^*) := \mathbf{e}^*(\mathbf{a})$. Then*

$$\mathbf{a} \in G^{(*)} \quad \text{if and only if} \quad ev_{\mathbf{a}} : T(V^*) \rightarrow \mathbb{R} \text{ is a morphism of algebras.} \quad (2.65)$$

The range of the signature is a proper subgroup of $G^{(*)}$, i.e. it is *not* the entirety of G^* (cf. Subsection 2.9). However, given any finite integer $n \in \mathbb{Z}_{\geq 0}$, the free nilpotent group $G^{(n)}$ of order n (cf. Definition 2.4) is the image of the truncated signature truncated to depth n of the space $\mathcal{V}^p([a, b], V)$. That is, $G^{(n)} = \Pi_n \circ S(\mathcal{V}^p([a, b], V))$ (see Proposition 2.27 in [Lyons et al., 2007]).

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 [Lyons et al., 2007] for details). If $x : [a, b] \rightarrow V$ is a path with signature $S_{a,b}(x) = \prod_{n=0}^{\infty} S_{a,b}^n(x)$ then we know that $S_{a,b}^0(x) \equiv 1$. Hence, if we define $\hat{S}_{a,b}(x) := S_{a,b}(x) - \mathbf{1} \in T((V))$, then we define the

log signature by taking the logarithm (defined in Definition 2.2) of the element $S_{a,b}(x) = \mathbf{1} + \hat{S}_{a,b}(x)$. This results in defining the Lie series $\text{LogSig}_{a,b}(x) \in \mathcal{L}((V))$ by (cf. (2.40))

$$\text{LogSig}_{a,b}(x) := \log S_{a,b}(x) \stackrel{(2.40)}{=} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n} \left(\hat{S}_{a,b}(x) \right)^{\otimes n}. \quad (2.66)$$

Recall that the signature $S_{a,b}(x) \in \tilde{T}((V))$ since $\pi_0(S_{a,b}(x)) = 1$. Thus, via Lemma 2.21 in [Lyons et al., 2007], taking the exponential (defined in Definition 2.2) 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 (2.56), given any positive integer $n \in \mathbb{Z}_{\geq 1}$ we define the *truncated log signature* (truncated to depth n) by $\text{LogSig}_{a,b}^{(n)}(x) := \Pi_n(\text{LogSig}_{a,b}(x))$. That is, the truncated log signature (truncated to depth n) is the projection of the full log signature to $\mathcal{L}^{(n)}(V)$, the Lie polynomials of degree n .

2.9. Rough Paths

In this subsection we outline how the theory of *rough paths* enables one to consider signatures in the setting of p -variation regularity for $p \geq 2$. The core idea is to view the signature as the core object rather than the path itself. In Subsection 2.8 we observed that if $x \in \mathcal{V}^p([a, b], V)$ for some $p \in [1, 2)$, then its signature $S(x)$ is a multiplicative functional $\Delta_{[a,b]} \rightarrow T((V))$ with finite p -variation in the sense that there exists a control $\omega : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ for which, given any $n \in \mathbb{Z}_{\geq 1}$ and any $(s, t) \in \Delta_{[a,b]}$, we have

$$\|S_{s,t}^n(x)\|_{V^{\otimes n}} \leq \frac{\omega(s, t)^n}{\Gamma(1 + n/p)}.$$

This motivates the following definitions.

Definition 2.6 (Multiplicative Functional). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, and $n \in \mathbb{Z}_{\geq 1}$. Then a continuous function $X : \Delta_{[a,b]} \rightarrow T^{(n)}(V)$ is called a *multiplicative functional of degree n* if it satisfies both of the following properties.

- (A) For every $(s, t) \in \Delta_{[a,b]}$ we have that $X_{s,t}^0 := \pi_0(X_{s,t}) = 1$.
- (B) For every $s, u, t \in [a, b]$ with $s \leq u \leq t$ we have that $X_{s,u} \otimes X_{u,t} = X_{s,t}$.

A continuous function $\Delta_{[a,b]} \rightarrow T((V))$ satisfying both properties (A) and (B) is referred to as a *multiplicative functional*.

Definition 2.7 (Finite p -Variation). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, $p \in \mathbb{R}_{\geq 1}$, and define $\beta \in \mathbb{R}_{\geq 1}$ by

$$\beta := p^2 \left(1 + \sum_{r=3}^{\infty} \left(\frac{2}{r-2} \right)^{\frac{|p|+1}{p}} \right). \quad (2.67)$$

Then a multiplicative functional $X : \Delta_{[a,b]} \rightarrow T((V))$ has *finite p -variation* if there exists a control $\omega : \Delta_{[a,b]} \rightarrow \mathbb{R}_{\geq 0}$ such that for every integer $n \in \mathbb{Z}_{\geq 1}$ and every $(s, t) \in \Delta_{[a,b]}$ we have

$$\|X_{s,t}^n\|_{V^{\otimes n}} \leq \frac{\omega(s, t)^{\frac{n}{p}}}{\beta \Gamma\left(1 + \frac{n}{p}\right)} \quad (2.68)$$

where Γ denotes the Gamma function.

The constant $\beta \in \mathbb{R}$ in (2.67) simplifies numerical constants in many of the fundamental results in the theory of rough paths. The precise value of β is not of particular importance; one could change its value without altering the class of multiplicative functionals determined by the requirement of finite p -variation.

We can now define a p -rough path for general $p \in \mathbb{R}_{\geq 1}$.

Definition 2.8 (p -Rough Path). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ be an interval, and $p \in \mathbb{R}_{\geq 1}$. Then a p -rough path in V is a multiplicative functional of degree $\lfloor p \rfloor$ (cf. Definition 2.6) having finite p -variation (cf. Definition 2.7). The collection of p -rough paths in V is denoted by $\Omega_p(V)$.

The space of p -rough paths $\Omega_p(V)$ is *not* closed under addition; having $X, Y \in \Omega_p(V)$ does not mean that $X + Y \in \Omega_p(V)$. Consequently $\Omega_p(V)$ is not a vector space.

Given a p -rough path $X \in \Omega_p(V)$, for $i \in \mathbb{Z}_{\geq 1}$ we call $X^i := \pi_i(X) \in V^{\otimes i}$ the i^{th} iterated integral of X despite there being no need for X^i to be related to any kind of integral.

The following *Extension Theorem* regarding p -rough paths is fundamental to the theory of rough paths.

Theorem 2.9 (Extension Theorem; Variant of Theorem 3.7 in [Lyons et al., 2007]). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, and $p \in \mathbb{R}_{\geq 1}$. Suppose that $X \in \Omega_p(V)$ is a p -rough path. Then there exists a unique multiplicative functional $\tilde{X} : \Delta_{[a,b]} \rightarrow T((V))$ with finite p -variation that coincides with X up to level $\lfloor p \rfloor$ in the sense that $\Pi_{\lfloor p \rfloor}(\tilde{X}) \equiv X$.

The Extension Theorem 2.9 allows us to make the following definition of the *signature* of a p -rough path $X \in \Omega_p(V)$.

Definition 2.10 (Signature of a p -rough path). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, $p \in \mathbb{R}_{\geq 1}$, and $X \in \Omega_p(V)$ be a p -rough path. The *signature* of X is the unique multiplicative functional $S(X) : [a, b] \rightarrow T((V))$ with finite p -variation such that $\Pi_{\lfloor p \rfloor}(S(X)) \equiv X$ in $T^{(\lfloor p \rfloor)}(V)$.

Once again we define the *log signature* of a p -rough path to be the Lie series in $\mathcal{L}((V))$ resulting from taking the logarithm defined in Definition 2.2 of the signature. If $X \in \Omega_p(V)$ is a p -rough path and $S(X) : [a, b] \rightarrow T((V))$ denote the signature of X , the log signature of X is the multiplicative functional $\text{LogSig}(X) : \Delta_{[a,b]} \rightarrow T((V))$ defined for $(s, t) \in \Delta_{[a,b]}$ by (cf. (2.40))

$$\text{LogSig}_{s,t}(X) := \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{n} (S(X) - \mathbf{1})^{\otimes n} \quad (2.69)$$

It is once again true that taking the exponential of the log signature recovers the signature. The log signature once again represents the same information as the signature but in a more compact form.

We observe that Definition 2.10 of the signature of a p -rough path is consistent with the definition of the signature given in Subsection 2.8 when $p \in [1, 2)$. Indeed if $p \in [1, 2)$ then a p -rough path $X \in \Omega_p(V)$ coincides with the signature of a path $x \in \mathcal{V}^p([a, b], V)$ as defined in Subsection 2.8. To be more precise, first suppose that $x \in \mathcal{V}^p([a, b], V)$. Then the signature $S(x)$ is the unique p -rough path $X \in \Omega_p(V)$ satisfying, for every $(s, t) \in \Delta_{[a,b]}$, that $\Pi_1(X_{s,t}) = (1, x_t - x_s) \in T^{(1)}(V)$. Conversely, let $X \in \Omega_p(V)$ for $p \in [1, 2)$. Then if we define $x \in \mathcal{V}^p([a, b], V)$ by defining $x_a := \pi_1(X_{a,a})$ and, for $t \in (a, b]$, $x_t := \pi_1(X_{a,t}) + x_a$, then it turns out that X coincides with the truncated signature $S^{(1)}(x)$ of this path x [Lyons et al., 2007]. The *Extension Theorem* 2.9 then ensures that $X \equiv S(x)$ as elements in the tensor algebra $T((V))$. In this case, given an integer $i \in \mathbb{Z}_{\geq 1}$, the projection $X^i := \pi_i(X)$ of X to $V^{\otimes i}$ is given by the i^{th} iterated integral of the path x , providing some justification for choosing to call X^i the i^{th} iterated integral of X .

An illustrative example is the so-called Brownian rough path. For this purpose we let $V = \mathbb{R}^d$ and consider an \mathbb{R}^d valued Brownian motion $B : [0, T] \rightarrow \mathbb{R}^d$ for some $T > 0$. For each $t \in [0, T]$ we will write $B_t = (B_t^1, \dots, B_t^d)$ as the decomposition of B_t with respect to the canonical basis e_1, \dots, e_d of \mathbb{R}^d . Then we may define a multiplicative functional $I : \Delta_{[0,T]} \rightarrow T^{(2)}(\mathbb{R}^d)$ by setting, for $(s, t) \in \Delta_{[0,T]}$,

$$I_{s,t} = \left(1, I_{s,t}^1, I_{s,t}^2\right) \in T^{(2)}(\mathbb{R}^d) \quad (2.70)$$

where

$$I_{s,t}^1 = \sum_{j=1}^d I_{s,t}^{1,j} e_j = \sum_{j=1}^d (B_t^j - B_s^j) e_j = B_t - B_s \quad (2.71)$$

and

$$I_{s,t}^2 = \sum_{i=1}^d \sum_{j=1}^d I_{s,t}^{2,ij} e_i \otimes e_j = \sum_{i=1}^d \sum_{j=1}^d \int_0^t \int_0^{u_2} dB_{u_1}^i dB_{u_2}^j e_i \otimes e_j. \quad (2.72)$$

The multiplicative functional I , which is sometimes called the Itô 2-multiplicative functional, is a p -rough path for any $2 < p < 3$ (see Corollary 3.17 in [Lyons et al., 2007]). The unique extension of I to a multiplicative functional $[0, T] \rightarrow T^{(2)}(\mathbb{R}^d)$ with finite p -variation then provides an interpretation for the signature of a Brownian motion.

Given the important role played by the notion of finite p -variation provided by Definition 2.7, we would like a notion of the p -variation distance between two p -rough paths. To do so, we first observe that $\Omega_p(V)$ is a subset of the vector space $C_p^0(\Delta_{[a,b]}; T^{([p])}(V))$ of continuous functions $\Delta_{[a,b]} \rightarrow T^{([p])}(V)$. We can equip $\Omega_p(V)$ with a p -variation metric d_p by first defining it on the linear space $C_p^0(\Delta_{[a,b]}; T^{([p])}(V))$. That is, for $X, Y \in C_p^0(\Delta_{[a,b]}; T^{([p])}(V))$ we define

$$d_p(X, Y) := \max_{i \in \{1, \dots, [p]\}} \sup_{a \leq t_0 \leq \dots \leq t_r \leq b} \left(\sum_{j=0}^{r-1} \left\| X_{t_j, t_{j+1}}^i - Y_{t_j, t_{j+1}}^i \right\|_{V^{\otimes i}}^{\frac{p}{i}} \right)^{\frac{1}{p}} \quad (2.73)$$

where for each $i \in \{1, \dots, [p]\}$ and every $(s, t) \in \Delta_{[a,b]}$ we have $X_{s,t}^i := \pi_i(X_{s,t})$ and $Y_{s,t}^i := \pi_i(Y_{s,t})$. If the space of p -rough paths is equipped with the p -variation metric then $(\Omega_p(V), d_p)$ is a complete metric space [Lyons and Qian, 2002]. A more convenient characterisation of the notion of convergence induced by d_p is provided in [Lyons and Qian, 2002] (see also Definition 3.12 in [Lyons et al., 2007]).

Whilst the notion of p -rough path defined in Definition 2.8 extends the theory from Subsection 2.8, it does so in an abstract theoretical manner. A more hands-on extension of the theory from Subsection 2.8 can be made as follows. Consider $q \in [1, 2)$ and a path $x \in \mathcal{V}^q([a, b], V)$. Then x defines a q -rough path $X \in \Omega_q(V)$ by setting $X_{s,t} := (1, x_t - x_s) \in T^{(1)}(V)$ for $(s, t) \in \Delta_{[a,b]}$. Then given any $p \geq q$ the *Extension Theorem* 2.9 tells us that this functional X can be extended to a multiplicative functional of degree $[p]$ with finite q -variation, and hence of finite p -variation.

In particular, a path with bounded variation (i.e. a path in $\mathcal{V}^1([a, b], V)$) can be canonically viewed as a p -rough path for any $p \geq 1$. Thus under this identification of an element $x \in \mathcal{V}^1([a, b], V)$ with a p -rough path $X \in \Omega_p(V)$, we can view $\mathcal{V}^1([a, b], V)$ as a subset of $\Omega_p(V)$. Hence we can consider the closure of $\mathcal{V}^1([a, b], V)$ in $\Omega_p(V)$ with respect to the metric d_p defined in (2.73).

Definition 2.11 (Geometric p -Rough Path). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, and $p \in \mathbb{R}_{\geq 1}$. Then a *geometric p -rough path* in V is a p -rough path $X \in \Omega_p(V)$ for which there exists a sequence of 1-rough paths $\{X(n)\}_{n=1}^\infty \subset \Omega_1(V)$ with $d_p(X(n), X) \rightarrow 0$ as $n \rightarrow \infty$. The space of geometric p -rough paths in V is denoted by $G\Omega_p(V)$.

It is evident that we have the inclusion that $G\Omega_p(V) \subset \Omega_p(V)$. However, the inclusion is strict; not all p -rough paths are geometric p -rough paths. For example, the Itô 2-multiplicative functional I defined in (2.70) is a p -rough path for $p \in (2, 3)$ but it is not geometric (see Section 3.3.1 in [Lyons et al., 2007]). But one can use an \mathbb{R}^d valued Brownian motion to determine a geometric p -rough path for $p \in (2, 3)$. This can be done by considering the dyadic piecewise linear approximation to the Brownian motion.

To be more precise, once again let $V = \mathbb{R}^d$ and consider an \mathbb{R}^d valued Brownian motion $B : [0, T] \rightarrow \mathbb{R}^d$ for some $T > 0$. For each $n \in \mathbb{Z}_{\geq 0}$ let $x(n) : [0, T] \rightarrow \mathbb{R}^d$ be the continuous path that is linear off the dyadic points $k/2^n$ and satisfies that $x(n)_{k/2^n} = B_{k/2^n}$. Further, for each $n \geq 0$ we can define a multiplicative functional $S(n) : \Delta_{[0,T]} \rightarrow T^{(2)}(\mathbb{R}^d)$ of degree 2 by setting, for each $(s, t) \in \Delta_{[0,T]}$,

$$S(n)_{s,t} := \left(1, \int_s^t dx(n)_u, \int_0^t \int_0^{u_2} dx(n)_{u_1} \otimes dx(n)_{u_2} \right) \in T^{(2)}(\mathbb{R}^d). \quad (2.74)$$

For any $p > 2$ the sequence $\{S(n)\}_{n=1}^\infty$ is a Cauchy sequence in the p -variation topology [Sipilainen, 1993,

[Friz, 2005]. Moreover its limit is given by

$$S_{s,t} = \left(1, B_t - B_s, \frac{1}{2}(B_t - B_s)^{\otimes 2} + A_{s,t} \right) \in T^{(2)}(\mathbb{R}^d) \quad (2.75)$$

where $A_{s,t}$ is the Lévy area of the Brownian motion. That is, for $i, j \in \{1, \dots, d\}$ we have

$$A_{s,t}^{ij} := \frac{1}{2} \int_0^T \int_0^{u_2} dB_{u_1}^i dB_{u_2}^j - dB_{u_1}^j dB_{u_2}^i. \quad (2.76)$$

The second iterated integral of the limit functional S is nothing but the second iterated Stratonovich integral of the Brownian motion; that is

$$S_{s,t}^2 \stackrel{(2.75)}{=} \frac{1}{2} (B_t - B_s)^{\otimes 2} + A_{s,t} = \int_0^T \int_0^{u_2} \circ dB_{u_1} \otimes \circ dB_{u_2}. \quad (2.77)$$

Consequently, the signature of the functional S is simply the Stratonovich signature of the Brownian motion, i.e. the multiplicative functional $X : \Delta_{[0,T]} \rightarrow T((\mathbb{R}^d))$ with, for $n \in \mathbb{Z}_{\geq 1}$, its n^{th} iterated integral $X^n : \Delta_{[0,T]} \rightarrow (\mathbb{R}^d)^{\otimes n}$ given by

$$X_{s,t}^n := \int_{s \leq u_1 \leq \dots \leq u_n \leq t} \circ dB_{u_1} \otimes \dots \otimes \circ dB_{u_n}. \quad (2.78)$$

This motivates the terminology that the multiplicative functional S defined in (2.75) is called the canonical Brownian rough path. Proof that $S \in G\Omega_p(\mathbb{R}^d)$ for $p \in (2, 3)$ may be found in [Sipilainen, 1993].

Recall from Subsection 2.8 that if $x \in \mathcal{V}^p([a, b], V)$ for a $p \in [1, 2)$ then, as an element in the tensor algebra $T((V))$, the signature $S(x)$ is group-like, i.e. $S(x) \in G^{(*)}$ where the group-like elements are defined in Definition 2.4. In particular, given an integer $n \in \mathbb{Z}_{\geq 0}$, the truncated signature $S^{(n)}(x)$ takes its values in the free nilpotent group $G^{(n)}$ of order n (cf. Definition 2.4). The algebraic properties related to the shuffle product that define $G^{(n)}$ ensure that a geometric p -rough path takes its values in $G^{(\lfloor p \rfloor)}$. Unfortunately this does *not* characterise geometric rough paths amongst all rough paths.

Definition 2.12 (Weakly Geometric p -Rough Path). Let V be a Banach space, $[a, b] \subset \mathbb{R}$ an interval, and $p \in \mathbb{R}_{\geq 1}$. Then a *weakly geometric p -rough path* in V is a p -rough path $X \in G\Omega_p(V)$ which takes its values in $G^{(\lfloor p \rfloor)}$, the free nilpotent group of order $\lfloor p \rfloor$. We denote the space of weakly geometric p -rough paths by $WG\Omega_p(V)$.

The collection of weakly geometric p -rough paths $WG\Omega_p(V)$ is distinct from both the collection of geometric p -rough paths $G\Omega_p(V)$, and the collection of p -rough paths $\Omega_p(V)$. Indeed we have the following strict inclusions:

$$G\Omega_p(V) \subset WG\Omega_p(V) \subset \Omega_p(V). \quad (2.79)$$

It is not immediately obvious that the first inclusion is strict; the interested reader may find the details in [Friz and Victoir, 2006]. The difference between geometric and weakly geometric rough paths is an annoying technicality that is, in spirit, similar to the distinction between C^1 and Lipschitz functions. The difference between geometric rough paths and general rough paths is, however, substantial. The class of weakly geometric rough paths is the most important one to bear in mind for all the applications considered in this article.

We end this subsection with a remark about the interaction of the p -variation topology and the q -variation topology for real numbers $q \geq p \geq 1$. Suppose that we have a p -rough path $X \in \Omega_p(V)$. Then the *Extension Theorem 2.9* tells us that X can be canonically extended to a q -rough path, and thus we can slightly abuse notation and write $X \in \Omega_q(V)$. In this way, any subset $\mathcal{K} \subset \Omega_p(V)$ can be viewed as a subset $\mathcal{K} \subset \Omega_q(V)$.

Now suppose that $\mathcal{K} \subset \Omega_p(V)$ is a compact subset, which, in particular, means that \mathcal{K} is a compact subset in the topology induced by the p -variation metric d_p defined in (2.73). Consider the map $E_{p,q} : \Omega_p(V) \rightarrow \Omega_q(V)$ mapping an element $Y \in \Omega_p(V)$ to its canonical extension in $\Omega_q(V)$. That is, $E_{p,q}(Y)$ is the truncation

to depth $\lfloor q \rfloor$ of the extension \tilde{Y} of Y provided by the *Extension Theorem 2.9*. It follows from Theorem 3.10 in [Lyons et al., 2007] that the map $E_{p,q}$ is continuous when the domain $\Omega_p(V)$ is equipped with the p -variation metric d_p defined in (2.73) and the target $\Omega_q(V)$ is equipped with the q -variation metric d_q defined in (2.73). Hence, being the continuous image of a compact set, the image $E_{p,q}(\mathcal{K})$ is a compact subset of $(\Omega_q(V), d_q)$.

Moreover, $E_{p,q}$ restricted to \mathcal{K} is an injective continuous mapping from a compact set \mathcal{K} into the Hausdorff space $(\Omega_q(V), d_q)$. Thus the restriction of $E_{p,q}$ to \mathcal{K} is in fact a homeomorphism onto its image, and so the q -variation topology on the image $E_{p,q}(\mathcal{K})$ coincides with the p -variation topology on \mathcal{K} .

Therefore, for every $q \geq p$ the image $E_{p,q}(\mathcal{K})$ is a compact subset of $(\Omega_q(V), d_q)$, and the q -variation topology induced on $E_{p,q}(\mathcal{K})$ by viewing it as a subset of $\Omega_q(V)$ coincides with the p -variation topology induced on \mathcal{K} as a subset of $(\Omega_p(V), d_p)$.

3. Signature as a Feature Map

We assume that V is a finite d -dimensional Banach space, and that $\Omega \subset \mathcal{S}(V) \times \mathbb{R}$ is a labelled dataset. Following our conventions, this means that each element in Ω is a pair $(\mathbf{x}, y) \in \mathcal{S}(V) \times \mathbb{R}$ where the stream \mathbf{x} is assumed to consist of incremental data. We assume that the streams of increments \mathbf{x} contained within the set Ω are of the form determined by the method of transforming a stream of values to a stream of increments detailed in Subsection 2.3. A particular consequence of this assumption is that the integer $d \in \mathbb{Z}_{\geq 1}$ must be even; see Subsection 2.3.

Suppose that $M \in \mathbb{Z}_{\geq 1}$ is finite, and that $\Omega = \{(\mathbf{x}_k, y_k) \in \mathcal{S}(V) \times \mathbb{R} \mid k \in \{1, \dots, M\}\}$. Recall that $\Omega_{\mathcal{S}(V)} \subset \mathcal{S}(V)$ denotes the projection of Ω onto $\mathcal{S}(V)$. Moreover, we assume that $\Omega_{\mathcal{S}(V)} \subset \mathcal{M} \subset \mathcal{S}(V)$, and our goal is to use the labelled dataset Ω to learn a continuous function $f \in C^0(\mathcal{M}; \mathbb{R})$ for the purpose of predicting the systems response to inputs $\mathbf{x} \in \mathcal{S}(V)$ that are not within $\Omega_{\mathcal{S}(V)}$. We reiterate that we are implicitly assuming that, in some sense, the subset \mathcal{M} is sufficiently well approximated by $\Omega_{\mathcal{S}(V)}$ to make this task reasonable (cf. Subsection 2.3).

We now outline a basic strategy for incorporating the signature as a feature map in this regression task.

Simple Signature Regression Framework

- (A) Consider the set of contours $C = \{\Gamma_{\mathbf{x}_k} \mid k \in \{1, \dots, M\}\}$ where, for each $k \in \{1, \dots, M\}$, the contour $\Gamma_{\mathbf{x}_k}$ is the contour formed by the concatenation of the entries within the stream of increments \mathbf{x}_k .
- (B) For each $k \in \{1, \dots, M\}$ choose a parameterisation $\gamma_k : [0, 1] \rightarrow V$ of the contour $\Gamma_{\mathbf{x}_k}$. For each $k \in \{1, \dots, M\}$, the stream \mathbf{x}_k contains a finite number of increments, thus each path γ_k is of bounded variation. That is, for each $k \in \{1, \dots, M\}$ we have $\gamma_k \in \mathcal{V}^1([0, 1], V)$, and so we can consider its signature $S(\gamma_k) : \Delta_{[0,1]} \rightarrow T((V))$.
- (C) Choose $K_0 \in \mathbb{Z}_{\geq 1}$ and consider the collection of truncated to depth K_0 signatures $S^{(K_0)}(\gamma_k)$ for $k \in \{1, \dots, M\}$. Numerically solve the system of equations given by

$$\sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{0,1}(\gamma_k)^{\mathbf{J}} = y_k \quad (3.1)$$

for each $k \in \{1, \dots, M\}$. This is a system of M equations in $1 + d + \dots + d^{K_0} = \frac{d^{K_0+1}-1}{d-1}$ variables.

- (D) Take the learnt function to be $f : \mathcal{M} \rightarrow \mathbb{R}$ defined by

$$f(\mathbf{v}) := \sum_{i=0}^{K_0} \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{0,1}(\gamma(\mathbf{v}))^{\mathbf{J}}, \quad (3.2)$$

where $\gamma(\mathbf{v}) : [0, 1] \rightarrow V$ is a parameterisation of the contour $\Gamma_{\mathbf{v}}$ corresponding to the stream $\mathbf{v} \in \mathcal{M}$.

The sensibility of this approach and the attractiveness of the signatures as a feature map from a machine learning point of view, as claimed in Section 2.6, rely on the following key theoretical properties of signatures. The first key property is that the infinite graded sequence of statistics provided by the signature identifies a contour in an essentially unique manner.

Theorem 3.1 (Uniqueness of signature, [Hambly and Lyons, 2010, Boedihardjo et al., 2016]). *Let V be a Banach space, $[a, b] \subset \mathbb{R}$ a compact interval, and $p \in \mathbb{R}_{\geq 1}$. Then the signatures of two geometric p -rough paths $X, Y \in G\Omega_p(V)$ coincide if and only if the X and Y are tree-like equivalent.*

The full technical details of *tree-like* equivalence may be found in either of [Hambly and Lyons, 2010] or [Boedihardjo et al., 2016], for example. Heuristically, the signature determines the path up to sections on which the path exactly retraces itself. In our setting such sections can be avoided by augmenting the initial streams to contain a monotonic channel. Recalling our view that the data is in the form of streams of increments, this can be achieved, for example, by augmenting each stream to have a channel consisting of strictly positive entries.

One particularly important example of tree-like equivalence is when one path is a translation of the other. To be more precise, suppose that $\gamma : [a, b] \rightarrow V$ is a continuous path of bounded variation. Consider a fixed $v \in V$ and define a continuous path of bounded variation $\psi : [a, b] \rightarrow V$ by setting, for $t \in [a, b]$, $\psi_t := \gamma_t + v$. Then γ and ψ are tree-like equivalent. Consequently the signature is invariant under translation of a path.

This invariance motivates our use of the invisibility reset transformation (cf. Subsection 2.5) when detailing our method of transforming a stream of values into a stream of increments in Subsection 2.3. Indeed the invisibility reset transformation is used to explicitly encode both the first and last values of a stream into its resulting stream of increments in a manner that is invariant under translation. Consequently, the signature of the path resulting from the concatenation of these increments *will* capture information related to the norm/size of the entries in the original stream of values. This information would *not* be captured by the signature if we worked with the stream of incremental changes obtained by only considering the differences between successive entries of the original stream of values.

Another important example of tree-like equivalence is when one path is a reparameterisation of the other. To be more precise, suppose that $\gamma : [a, b] \rightarrow V$ is a continuous path of bounded variation. Then if $\psi : [a, b] \rightarrow [a, b]$ is a continuously differentiable increasing surjection, then the paths γ and $\gamma \circ \psi$ are tree-like equivalent. Consequently the signature is invariant under reparameterisation of a path; it essentially quotients out the infinite-dimensional group of time reparameterisation. This formalises our earlier claims that the signature filters out the parameterisation noise.

A consequence of this reparameterisation invariance is that the signature can record the order of events *without* needing to record precisely when each event occurs. This is a distinct 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 [Levin et al., 2013].

Theorem 3.1 ensures that, for each $k \in \{1, \dots, M\}$, the choice of parameterisation γ_k for the contour Γ_k is irrelevant to the resulting signature $S(\gamma_k)$. Any two parameterisations of Γ_k are reparameterisations of one and other, and hence Theorem 3.1 means their signatures coincide. In fact this provides a way to determining a topology on a space of unparameterised paths in V .

Indeed, given a fixed compact interval $[a, b] \subset \mathbb{R}$, we may consider the space of unparameterised paths in V determined by taking the quotient of $\mathcal{V}^1([a, b], V)$ with respect to the equivalence relation generated by the notion of tree-like equivalence [Hambly and Lyons, 2010]. We will denote the resulting space by $\text{Path}_1([a, b], V)$, and each element within it is an equivalence class $[\gamma]$ for some $\gamma \in \mathcal{V}^1([a, b], V)$. By appealing to Theorem 3.1, we can define the signature of an element $[\gamma] \in \text{Path}_1([a, b], V)$ to be the signature of *any* representative $\alpha \in \mathcal{V}^1([a, b], V)$ of the equivalence class $[\gamma]$.

There is no canonical topology on $\text{Path}_1([a, b], V)$ [Hambly and Lyons, 2010]. A thorough investigation of various approaches to equipping $\text{Path}_1([a, b], V)$ with a topology and the subsequent properties

of each resulting topology is provided in [Cass and Turner, 2024]. A concrete way to use the signature to determine a topology on $\text{Path}_1([a, b], V)$ is to define a metric on $\text{Path}_1([a, b], V)$ as follows. Given $[\gamma], [\alpha] \in \text{Path}_1([a, b], V)$ define

$$d_{\text{Path}_1}([\gamma], [\alpha]) := d_1(S(\gamma), S(\alpha)) \quad (3.3)$$

where $S(\gamma), S(\alpha) : \Delta_{[a, b]} \rightarrow T((V))$ denote the signatures of γ and α respectively, and d_1 is the 1-variation metric defined in (2.73). Theorem 3.1 ensures that d_{Path_1} defined in (3.3) is well-defined.

The next key property is that the terms of the signature decay in size factorially as the order increases.

Theorem 3.2 (Factorial Decay; Proposition 2.2 in [Lyons et al., 2007]). *Let V be a Banach space, $[a, b] \subset \mathbb{R}$ a compact interval, and $\gamma \in \mathcal{V}^1([a, b], V)$. Let $S(\gamma) : \Delta_{[a, b]} \rightarrow T((V))$ denote the signature of γ . Then for every $n \in \mathbb{Z}_{\geq 1}$ and every $(s, t) \in \Delta_{[a, b]}$, the n^{th} component of the signature $S^n(\gamma) : \Delta_{[a, b]} \rightarrow V^{\otimes n}$ satisfies that*

$$\|S_{s, t}^n(\gamma)\|_{V^{\otimes n}} \leq \frac{\|\gamma\|_{1, [a, b]}^n}{n!}. \quad (3.4)$$

Consequently, in a norm based sense, the truncation of the signature to a finite depth manages to capture the terms with the most significant contribution. Thus it is reasonable, again in a norm based sense, to imagine that by choosing a sufficiently large K_0 in **Simple Signature Regression Framework** step (C), the resulting truncated signatures still manage to summarise the most significant affects of the paths.

Examining **Simple Signature Regression Framework** steps (C) and (D), the underlying idea is to approximate the systems response by a linear combination of signature components. The sensibility of this relies upon which classes of functions are captured by such linear combinations. Loosely, we need to know if linear combinations of signature components are capable of providing a good approximation to an arbitrary continuous function. The final key property is that this is the case since continuous functions of paths are approximately linear on signatures. Thus the signature is, in some sense, a *universal nonlinearity* on paths.

Theorem 3.3 (Universal Nonlinearity; Variant of Theorem 2.1 in [Lemerrier et al., 2021]). *Let V be a Banach space, $[a, b] \subset \mathbb{R}$ a compact interval, $\mathcal{K} \subset \text{Path}_1([a, b], V)$ a compact subset, and $F : \mathcal{K} \rightarrow \mathbb{R}$ a continuous function. Then for any $\varepsilon > 0$ there exists a truncation level $n \in \mathbb{Z}_{\geq 1}$ and a collection of real coefficients $\{\alpha_i(\mathbf{J}) \mid i \in \{0, \dots, n\} \text{ and } \mathbf{J} \in \{1, \dots, d\}^i\}$ such that for every $[\theta] \in \mathcal{K}$ we have*

$$\left| F([\theta]) - \sum_{i=0}^n \sum_{\mathbf{J} \in \{1, \dots, d\}^i} \alpha_i(\mathbf{J}) S_{a, b}(\theta)^{\mathbf{J}} \right| \leq \varepsilon. \quad (3.5)$$

In Theorem 3.3 the topology on $\text{Path}_1([a, b], V)$ is the topology determined by the metric d_{Path_1} defined in (3.3). We sketch one approach to establishing this result. The image $\{S(\gamma) \mid [\gamma] \in \mathcal{K}\}$ determines a compact subset of $\Omega_1(V)$. Further, the shuffle product relation (2.62) can be used to establish that the coordinate iterated integrals span an algebra. Then the conclusion of Theorem 3.3 is obtained by appealing to the *Stone–Weierstrass Theorem* [Stone, 1948].

Returning to **Simple Signature Regression Framework** step (A), the subset $\{\Gamma_k \mid k \in \{1, \dots, M\}\}$ determines a finite subset of $\text{Path}_1([0, 1], V)$. Hence Theorem 3.3 is applicable, and we are guaranteed that, provided a sufficiently large truncation level is chosen, it will be possible to approximate *any* continuous function on this finite subset by a linear combination of the coordinate iterated integrals of the truncated signature.

We have only examined the signatures of the contours corresponding to the data set at the 1-variation scale. An obvious motivation for doing so is that the contours naturally make sense at this scale; that is, as paths with bounded variation. But just because they make sense at this scale does not mean it is the most natural scale to work with. It might be more sensible to consider them at the p -variation scale for $p > 1$.

As observed in Subsection 2.9, if $\gamma \in \mathcal{V}^1([a, b], V)$ and $p \geq 1$ then the truncation of the signature of γ to depth $[p]$ give a canonical way to lift γ to a p -rough path $X(p) = S^{([p])}(\gamma) \in \Omega_p(V)$. Consequently one can canonically view $\mathcal{V}^1([a, b], V)$ as a subset of $\Omega_p(V)$, and choose to work with the p -variation topology induced by the p -variation distance metric defined in 2.73. For each $j \in \{1, \dots, M\}$ and any $p \in \mathbb{R}_{\geq 1}$ let $X_j(p)$ denote the canonical lift of γ_j to a p -rough path in $\Omega_p(V)$, i.e. $X_j(p) := S^{([p])}(\gamma_j)$.

Then as remarked at the end of Subsection 2.9, the subset $\{X_1(p), \dots, X_M(p)\}$ is both compact as a subset in $(\Omega_p(V), d_p)$, and additionally satisfies that the p -variation topology induced on $\{X_1(p), \dots, X_M(p)\}$ coincides with the 1-variation topology induced on $\{X_1(1), \dots, X_M(1)\}$ as a subset of $(\Omega_1(V), d_1)$. Consequently, the theoretical guarantees of Theorem 3.3 remain valid even after switching to viewing our collection of signatures at the p -variation scale for some $p > 1$.

The central idea has been to associate a particular p -rough path with each stream \mathbf{x}_k in the original data set $\Omega_{S(V)}$. Thus far we have considered doing so by first associating a 1-rough path with the stream \mathbf{x}_k , and, if deemed sensible, subsequently lifting this 1-rough path to a p -rough path for $p > 1$. The incremental information provided by the stream \mathbf{x}_k determines how the lift is defined. However, if the stream \mathbf{x}_k contains more than just incremental data, there may be a more natural way of prescribing the higher order information in order to associate \mathbf{x}_k with a p -rough path.

For example, suppose a stream provided increments corresponding to an \mathbb{R}^d valued Brownian motion $B : [0, T] \rightarrow \mathbb{R}^d$ for some $T > 0$. Then both the Itô 2-multiplicative functional (cf. (2.70)) and the canonical Brownian rough path (cf. (2.75)) defined in Subsection 2.9 determine distinct p -rough paths, for $p \in (2, 3)$, which could be associated to the stream. However, if the stream additionally contained information determining how, for example, the area evolves, then which of these two candidates is more appropriate to choose could be determined.

The point is that we could associate a p -rough path to a stream \mathbf{x} by appropriately prescribing the coordinate iterated integrals up to depth $\lfloor p \rfloor$. The *Extension Theorem* 2.9 would then tell us that all the coordinate iterated integrals at higher depths would be uniquely determined, and so we could follow the method outlined in **Simple Signature Regression Framework** using the signatures of these prescribed p -rough paths. Consideration of Theorem 3.1 suggests it is most sensible to associate each stream with a geometric p -rough path, since we then preserve the same sense of uniqueness for the signature.

Indeed with this tree-like equivalence being preserved we could repeat our earlier approach. We can define $\text{Path}_p([a, b], V)$ to be the quotient of $G\Omega_p(V)$ with respect to the equivalence relation generated by the notion of tree-like equivalence. Once again Theorem 3.1 ensures we can define the signature of an element $[X] \in \text{Path}_p([a, b], V)$ to be the signature of the geometric p -rough path $Y \in G\Omega_p(V)$ for *any* representative Y of the equivalence class $[X]$. The p -variation metric d_p on $\Omega_p(V)$ can now be used to define a metric on $\text{Path}_p([a, b], V)$. Given $[X], [Y] \in \text{Path}_p([a, b], V)$ define

$$d_{\text{Path}_p}([X], [Y]) := d_p(S(X), S(Y)) \quad (3.6)$$

where $S(X), S(Y) : \Delta_{[a, b]} \rightarrow T((V))$ denote the signatures of X and Y respectively, and d_p is the p -variation metric defined in (2.73). Theorem 3.1 ensures that d_{Path_p} defined in (3.3) is well-defined.

The analogue of Theorem 3.3 in this setting can then be established in exactly the same way. Multiplicativity ensures the shuffle product remains valid, and so the coordinate iterated integrals (i.e. the components of the signature) still span an algebra. Hence the *Stone–Weierstrass* theorem [Stone, 1948] may once again be invoked to determine that on compact subsets of $\text{Path}_p([a, b], V)$ continuous functions can be arbitrarily well-approximated by linear combinations of coordinate iterated integrals (i.e. signature components).

The theoretical guarantees for the strategy outlined in **Simple Signature Regression Framework** make the signature an attractive choice of feature map. The signature summarises a path in a way that is essentially unique and, in particular, that is invariant under reparameterisation of the path (cf. Theorem 3.1). The coefficients of the signature (namely, the coordinate iterated integrals) provide a set of feature functions that have a natural grading (cf. Theorem 3.2), and which are rich enough to capture continuous functions on path space (cf. Theorem 3.3).

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 3.3). 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* [Lyons, 1998, Lyons et al., 2007, Friz and Victoir, 2010], 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 [Lyons et al., 2007]. A comprehensive exposition is provided in the textbook [Friz and Victoir, 2010]. CDEs have been used within the field of *Deterministic Control Theory*; a comprehensive introduction to deterministic control theory can be found in [Zabczyk, 2020].

The link with signatures is apparent from our earlier observation in Section 2.7 (cf. (2.38)) 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

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

for every $t \in (t_0, t_1]$. 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

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

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 3.2 in this article); see [Lyons et al., 2007, Lyons, 2014], for example. Existence and uniqueness results for the general case using *rough path theory* may be found, for example, in [Lyons, 1998, Lyons et al., 2007, Friz and Victoir, 2010].

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 [Lyons, 2014]. The link between truncated log signatures and vector fields (see Section 6 in [Lyons, 2014], 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 numeric 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 [Morrill et al., 2021b] and their subsequent refinement

in [Walker et al., 2024]. In [Foster et al., 2020] 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 [Foster et al., 2020] 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 [Morrill et al., 2021b] 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 [Boutaib et al., 2013], 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 $f^{\circ(k)} : W \rightarrow \mathbf{L}(V^{\otimes k}, W)$ using the derivatives of f as follows (see Definition A.6 in [Morrill et al., 2021b]). Below we use the notation that $Df^{\circ(k-1)} : W \rightarrow \mathbf{L}(W, \mathbf{L}(V^{\otimes(k-1)}, W))$ denotes the Fréchet derivative of $f^{\circ(k-1)}$. 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.

- For $k = 0$ we define $f^{\circ(0)} : W \rightarrow W$ by setting $f^{\circ(0)}(y) := y$ for every $y \in W$.
- For $k = 1$ we define $f^{\circ(1)} : W \rightarrow \mathbf{L}(V, W)$ by setting $f^{\circ(1)}(y) := f(y)$ for every $y \in W$.
- For the remaining $k \in \{2, \dots, N\}$ we proceed inductively. Once $f^{\circ(k-1)} : W \rightarrow \mathbf{L}(V^{\otimes(k-1)}, W)$ has been defined, define $f^{\circ(k)} : W \rightarrow \mathbf{L}(V^{\otimes k}, W)$ as follows. For each $y \in W$ we set $f^{\circ(k)}(y)$ to be the unique k -linear map in $\mathbf{L}(V^{\otimes k}, W)$ for which

$$f^{\circ(k)}(y)[v \otimes v_{k-1}] = Df^{\circ(k-1)}(y)(f(y)[v])[v_{k-1}] \quad (4.4)$$

whenever $v \in V$ and $v_{k-1} \in V^{\otimes(k-1)}$.

As noted in [Morrill et al., 2021b] the functions $f^{\circ(0)}, \dots, f^{\circ(N)}$ naturally arise in the Taylor expansion associated to the CDE (4.1). A combination of the functions $f^{\circ(0)}, \dots, f^{\circ(N)}$ 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 computed 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

$$\tilde{f}(u) := \sum_{k=0}^N f^{\circ(k)}(u) \left[\pi_k \left(\text{LogSig}_{s,t}^{(N)}(X) \right) \right] \quad (4.5)$$

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

$$x(0) = z_s \quad \text{and} \quad \frac{dx(r)}{dr} = \tilde{f}(x(r)) \quad (4.6)$$

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

$$\text{LogODE}(z_s, f, \text{LogSig}_{s,t}^{(N)}(X)) := x(1). \quad (4.7)$$

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

$$\text{LogODE}(z_s, f, \text{LogSig}_{s,t}^{(N)}(X), \tau) := x\left(\frac{\tau - s}{t - s}\right) \quad (4.8)$$

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 (4.5) (see Remark A.10 in [Morrill et al., 2021b]). This is a consequence of the fact that the log signature of X lies in a certain free Lie algebra;

see Subsections 2.7, 2.8, and 2.9 in this article, or refer Section 2.2.4 of [Lyons et al., 2007] for the precise details.

Error estimates for the log-ODE method may be found in [Boutaib et al., 2013] (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 [Morrill et al., 2021b] 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 [Lyons et al., 2007, Friz and Victoir, 2010], 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 [Boutaib et al., 2013].

The more recent work [Bayer et al., 2023a] develops an adaptive algorithm for effective use of the log-ODE method for solving *Rough Differential Equations* RDEs. This method involves a problem-dependent way of determining how to choose the two hyper-parameters governing the log-ODE method; namely, the length of the time intervals to consider the approximating ODE over, and the depth of log signature to consider. There is no universal right answer for choosing these parameters; the appropriate choice will depend on the differential equation through the regularity of the vector field f , the complexity of the paths, and the level of accuracy one desires.

The error representation formula derived by the authors in [Bayer et al., 2023a] is the main tool for understanding the relationship between these three aspects for the log-ODE method. In particular, improved a posteriori uniform error estimates are achieved for a class of RDE approximation schemes in Section 3 of [Bayer et al., 2023a], and it is further established that this class includes the log-ODE method in Section 4 of [Bayer et al., 2023a].

In a very loose sense these estimates can be summarised as follows. Provided the vector field f is smooth enough, that the time interval over which the approximation is sought is short enough, and that the truncation level for the log signature is deep enough, the error between the actual solution and the approximation resulting from the log-ODE method behaves like a power of the p -variation norm of the path X . Quantifying this statement is a far from trivial task which we make no effort to do; the reader seeking the rigorous detailed statements is directed to [Bayer et al., 2023a].

We end this section with the following remarks. It follows from the error estimates of [Bayer et al., 2023a] that one can always guarantee arbitrarily strong estimates for the log-ODE method by working over sufficiently small time intervals (i.e. by making $t - s$ sufficiently small). But it is *not* true that on a fixed time interval (i.e. a fixed value of $t - s$) that one can obtain arbitrarily strong estimates for the log-ODE method by choosing a sufficiently high truncation order for the log signature.

5. Computing Signatures

Making use of rough path and signature based techniques only becomes feasible if one is able to explicitly compute signatures of paths. Software packages for managing such computations have evolved over time. Initially there were python packages for computing path signatures directly from time series represented as NumPy arrays of doubles, see either *esig* [Lyons,] originating in 2002, or *iisignature* [Reizenstein and Graham, 2018]. These packages could compute signatures and log signatures, but did not support efficient algorithms for computing signatures over multiple intervals, or any ability to manipulate the signature or log signature tensors, e.g. exponentiate the log signature, or apply the Campbell–Baker–Hausdorff product.

Esig and *iisignature* are tools intended to be used without GPU based acceleration. Differentiable and GPU based tools with similar functionality, suitable for deep learning, include *Signatory* [Kidger and Lyons, 2020] and *signax* [Tong,]; *signatory* depends upon *PyTorch* [Paszke et al., 2019] while *signax* depends upon *JAX* [Bradbury et al., 2018].

A C++ package *libalgebra*, which has been in development since 2002 and remains in very active development, is a powerful and comprehensive CPU based library with tools for manipulating signatures, Lie elements, tensor algebra elements, and the corresponding dual objects (there are implementations of shuffle and half shuffle). This templated code has been optimised for floating point coefficients, but importantly also supports rational coefficients, arbitrary fixed precision coefficients, and even coefficients that are polynomials.

The former allows one to experiment and identify the precision required to execute high order calculations effectively. The latter allows one to develop formulae for complex Lie operations, and is also extremely valuable for debugging when code does not produce expected outcomes.

Esig wrapped some of the most elementary functionality from an early version of libalgebra.

RoughPy [Morley and Lyons, 2024] is a major upgrade on the python interface to managing rough streams. It allows the construction of a stream object from raw data; one can then query this object over *any* interval to get the signature or log signature of the stream over the specified interval. It is efficient and substantially streamlines the analysis of streamed data. RoughPy has all the functionality of the C++ package libalgebra mentioned above. In addition, it has automated GPU acceleration through **OpenCL** support. It wraps a new C++ library **libalgebra-lite** that is based upon libalgebra but with a reduced use of templates to allow better interaction with the python wrapper. In this survey we use RoughPy. The documentation for RoughPy may be found at roughpy.org. RoughPy may be installed directly via **pip** (a standard package-management system used to install and manage software packages written in Python).

In the remainder of this section we give a brief introduction to some of the functionality of RoughPy. A more in-depth introduction to RoughPy can be found in [Morley and Lyons, 2024], for example.

The main object in RoughPy is the **stream** object. The data is stored as a stream and, for example, signatures and log signatures are computed by querying the stream over an **interval** chosen by the user. The convention in RoughPy is that intervals are taken to be half-open in the sense that the start point is included but the end point is not. To be more precise, the python code `rp.RealInterval(0.0,2.0)` would generate the interval $[0, 2) \subset \mathbb{R}$.

The convention of using half-open intervals is just one aspect of the careful way in which RoughPy works with intervals. The decision to take half-open intervals is done to ensure that RoughPy can handle jumps in the underlying signal that happen at *any* particular time whilst still guaranteeing that combining signatures over adjacent intervals returns the signature over the entire interval determined by their union.

Moreover, for efficiency, an internal representation of a stream involving caching the signature over dyadic intervals of different resolutions is used. Resolution refers to the length of the finest granularity at which information about the underlying data is stored. Using the cache to recover the signature over any interval has logarithmic complexity. If n is the internal resolution of the stream, then recovering the signature over any interval uses, at most, $2n$ tensor multiplications. Any event in the stream occurs within one of these finest granularity intervals. If multiple events occur within the same interval, RoughPy resolves to a more complex log-signature which correctly reflects the time sequence of the events within the interval.

No query of the stream can see finer resolution than the internal resolution of the stream; a query can only access information over intervals that are a union of these finest resolution granular intervals. Hence a query over any interval is replaced by a query over an interval whose endpoints are shifted to be consistent with the granular resolution. This is obtained by rounding each endpoint to the contained end-point of the unique half-open granular interval containing this point. In particular, if both the left-hand and right-hand ends of the interval are contained in the half-open granular interval, the interval is rounded to the empty interval. Specifying a resolution of 32 or 64 equates to using integer arithmetic.

We first illustrate the use of 'LieIncrementStream' for the purpose of computing a truncated signature of a given stream of increments. Suppose $N, d \in \mathbb{Z}_{\geq 1}$ and we have a stream of increments $\mathbf{x} \in \mathcal{S}(\mathbb{R}^d)$ given by $\mathbf{x} = (I_1, \dots, I_N)$ for $I_1, \dots, I_N \in \mathbb{R}^d$. Then the python code in Figure 1 computes the signature of \mathbf{x} truncated to depth $k \in \mathbb{Z}_{\geq 1}$.

```
>>> import numpy as np
>>> import roughpy as rp
>>> x = np.array([ I_1, ..., I_N ])
>>> stream = rp.LieIncrementStream.from_increments(x , depth = k)
>>> interval = rp.RealInterval(0,N)
>>> sig = stream.signature(interval, depth=k)
```

Figure 1: Python code to compute the signature of the stream \mathbf{x} truncated to depth k using RoughPy

In Figure 1 'stream' is a stream whose (hidden) underlying data are the N increments $I_1, \dots, I_N \in \mathbb{R}^d$ and whose algebra elements are truncated to a maximum depth k . By default the increments are assumed to

occur at parameter values equal to their row index in the provided data. Thus here I_1 is assumed to occur at 0, I_2 is assumed to occur at 1, and so on until I_N is assumed to occur at $N - 1$. Consequently, all parameters are contained within the interval $[0, N)$ that is considered in Figure 1 via the code ‘interval = rp.RealInterval(0,N)’.

Instead of manually providing a depth when creating the stream ‘stream’, we could provide a **context**. A context is used to tell RoughPy the desired shape that streams should be assumed to have. Thus the same output achieved in Figure 1 could be obtained using the python code in Figure 2.

```
>>> import numpy as np
>>> import roughpy as rp
>>> x = np.array([ I_1, ..., I_N ])
>>> context = rp.get_context(width = d, depth = k, coeffs = rp.DPReal)
>>> stream = rp.LieIncrementStream.from_increments(x, ctx = context)
>>> interval = rp.RealInterval(0,N)
>>> sig = stream.signature(interval, ctx = context)
```

Figure 2: Python code to compute the signature of the stream \mathbf{x} truncated to depth k using RoughPy via a prescribed context.

The ‘context’ defined in Figure 2 has width d (since each entry in the stream \mathbf{x} is an element in \mathbb{R}^d), depth k (since we want to truncate to depth k), and real coefficients (since we are working over the reals).

The output ‘sig’ in both Figures 1 and 2 is a **FreeTensor** of width d , depth k , and ctype ‘DPReal’. Moreover, when printed, ‘sig’ returns both the values of the coefficients of the truncated signature and the basis element to which each coefficient corresponds. We illustrate this via the following explicit example. Fix $d = 2$, $k = 3$, and consider the stream of increments $\mathbf{x} = ((1, 1)) \in \mathcal{S}(\mathbb{R}^2)$. Then we compute the signature of this stream \mathbf{x} truncated to depth 3 using the python code in Figure 3.

```
>>> import numpy as np
>>> import roughpy as rp
>>> x = np.array([ [1.0,1.0] ])
>>> context = rp.get_context(width=2,depth=3, coeffs = rp.DPReal)
>>> stream = rp.LieIncrementStream.from_increments(x, ctx = context)
>>> interval = rp.RealInterval(0,1)
>>> sig = stream.signature(interval, ctx = context)
>>> sig
FreeTensor(width=2, depth=3, ctype=DPReal)
>>> print(sig)
{ 1() 1(1) 1(2) 0.5(1,1) 0.5(1,2) 0.5(2,1) 0.5(2,2) 0.166667(1,1,1)
0.166667(1,1,2) 0.166667(1,2,1) 0.166667(1,2,2) 0.166667(2,1,1)
0.166667(2,1,2) 0.166667(2,2,1) 0.166667(2,2,2) }
```

Figure 3: Python code to compute signature of stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^2)$ truncated to depth 3. Including the returned output of the commands ‘sig’ and ‘print(sig)’.

As expected, ‘sig’ is a FreeTensor of width 2, depth 3, and ctype DPReal. The output returned by the command ‘print(sig)’ is the collection of coefficients determining the signature truncated to depth 3 along with the word (in brackets) determining which coordinate iterated integral each coefficient represents. For example, the entry 0.5(1, 2) means that the value of the coordinate iterated integral corresponding to the word (1, 2) is 0.5. We recall from Subsection 2.8 that in this simple setting this is just the value of the iterated integral $\int_0^1 \int_0^y dx dy$ where we have associated the canonical x and y coordinates of \mathbb{R}^2 with ‘1’ and ‘2’ respectively.

If one was instead interested in the log signature of the stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^2)$, inserting the line of code

```
>>> logsigA = stream.log_signature(interval,ctx=context) (5.1)
```

into the code presented in Figure 3 would return the log signature truncated to depth 3. We could also compute

the log signature by directly taking the logarithm of the already computed signature. That is, we could use the python code

```
>>> logsigB = sig.log() (5.2)
```

to compute the desired log signature. However, there is a distinction between the outputs `logsigA` and `logsigB` of the approaches (5.1) and (5.2) respectively. The respective outputs are illustrated in Figure 4.

```
>>> logsigA
Lie(width=2, depth=3, ctype=DPReal)
>>> print(logsigA)
{ 1(1) 1(2) }
>>> logsigB
FreeTensor(width=2, depth=3, ctype=DPReal)
>>> print(logsigB)
{ 1(1) 1(2) -2.77556e-17(1,1,1) -2.77556e-17(1,1,2)
-2.77556e-17(1,2,1) -2.77556e-17(1,2,2) -2.77556e-17(2,1,1)
-2.77556e-17(2,1,2) -2.77556e-17(2,2,1) -2.77556e-17(2,2,2) }
```

Figure 4: The outputs resulting from the commands '`logsigA`', '`print(logsigA)`', '`logsigB`' and '`print(logsigB)`' where `logsigA` and `logsigB` are the results of the python code in (5.1) and (5.2) respectively.

The first difference is that `logsigA` is a **Lie** element whilst `logsigB` is a `FreeTensor`. The essential difference is that a Lie element is in the Free Lie algebra generated by '1' and '2' and recorded with respect to a *Hall basis*, whilst a `FreeTensor` is in the Tensor algebra (cf. Subsection 2.7) recorded with respect to the basis generated by all string combinations of '1' and '2'.

The second difference is a result of numerical errors being introduced due to truncation. Instead of taking log of the full signature tensor, we are only taking log of the signature truncated to depth 3 to compute `logsigB`. Consequently numerical errors enter the computations. In particular, the non-zero coefficients for the basis elements involving three entries from $\{1, 2\}$ in Figure 4, all of magnitude 10^{-17} , should all be zero.

We can also go from log signatures to signatures by exponentiation of a log signature. However, in order to do so we must first convert the log signature from a Lie element to a `FreeTensor`, and then exponentiate. As an example, consider a stream of increments $\mathbf{x} = ((1, 0, 0), (0, 1, 0), (0, 0, 1)) \in \mathcal{S}(\mathbb{R}^3)$. Then the following python code computes the signature of \mathbf{x} truncated to depth 2 both directly and by taking the exponential of the log signature.

As illustrated in Figure 5, both methods of computing the signature result in the same output for this example. It is worth remarking that the '`lie_to_tensor`' function requires a context to inform it about the shape of the input and output.

We can now provide the full illustration of the signatures ability to record the order of events as commented on in Subsection 2.2. For this purpose we recall the toy problem considered in Subsection 2.2. Let $N \in \mathbb{Z}_{\geq 1}$ and suppose we have a finite collection of time series, with each time series consisting of 2 channels and N samples. Suppose that each time series \mathbf{x} has the following structure. If $\mathbf{x} = \{(x_{1,i}, x_{2,i})\}_{i=1}^N$ then the following properties are true.

- For every $i \in \{1, \dots, N\}$ we have $x_{1,i}, x_{2,i} \in \{0, 1\}$.
- For $j \in \{1, 2\}$ if $i \in \{1, \dots, N\}$ and $x_{j,i} = 1$, then for every $k \in \{i, \dots, N\}$ we have $x_{j,k} = 1$.
- At least one of $x_{1,N}$ and $x_{2,N}$ is equal to 1.

Consider the task of determining which channel is the first to change from 0 to 1.

Recall from Subsection 2.5 that viewing each time series \mathbf{x} as an element in \mathbb{R}^{2N} and trying to learn a quadratic polynomial capable of determining which channel is the first to change from 0 to 1 involves considering $1 + 3N + 2N^2$ linearly independent quadratic polynomials. If we consider $N = 3600$ then we end up dealing with 25930801 linearly independent quadratic polynomials.

But computing the depth 2 signature will determine which channel changes first. Observe that there are only two possible options. Either the stream values go in the order $(0, 0) \rightarrow (1, 0) \rightarrow (1, 1)$ or the

```

>>> import numpy as np
>>> import roughpy as rp
>>> x = np.array([ [1.0,0.0,0.0] , [0.0,1.0,0.0] , [0.0,0.0,1.0] ])
>>> context = rp.get_context(width=3,depth=2, coeffs = rp.DPReal)
>>> stream = rp.LieIncrementStream. from_increments(x, ctx = context)
>>> interval = rp.RealInterval(0,1)
>>> sigA = stream.signature(interval, ctx = context)
>>> sigA
FreeTensor(width=3, depth=2, ctype=DPReal)
>>> print(sigA)
{ 1(0) 1(1) 1(2) 1(3) 0.5(1,1) 1(1,2) 1(1,3) 0.5(2,2) 1(2,3) 0.5(3,3) }
>>> logsig = stream.log_signature(interval, ctx = context)
>>> logsigT = context.lie_to_tensor(logsig)
>>> sigB = logsigT.exp()
>>> sigB
FreeTensor(width=3, depth=2, ctype=DPReal)
>>> print(sigB)
{ 1(0) 1(1) 1(2) 1(3) 0.5(1,1) 1(1,2) 1(1,3) 0.5(2,2) 1(2,3) 0.5(3,3) }

```

Figure 5: Python code to compute signature of stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^3)$ truncated to depth 2 both directly and by computing the exponential of the log signature.

order $(0,0) \rightarrow (0,1) \rightarrow (1,1)$. Hence the channels of the increments are either equivalent to the stream $\mathbf{x}_1 := ((1,0), (0,1)) \in \mathcal{S}(\mathbb{R}^2)$ or to the stream $\mathbf{x}_2 := ((0,1), (1,0)) \in \mathcal{S}(\mathbb{R}^2)$. The RoughPy output when the signature of \mathbf{x}_1 truncated to depth 2 is computed is

$$\{1(0) \quad 1(1) \quad 1(2) \quad 0.5(1,1) \quad 1(1,2) \quad 0.5(2,2)\}. \quad (5.3)$$

The RoughPy output when the signature of \mathbf{x}_1 truncated to depth 2 is computed is

$$\{1(0) \quad 1(1) \quad 1(2) \quad 0.5(1,1) \quad 1(2,1) \quad 0.5(2,2)\}. \quad (5.4)$$

We immediately see that the channel that is first to change from 0 to 1 is captured by the depth 2 coordinate iterated integrals corresponding to the words $(1,2)$ and $(2,1)$. If the x -coordinate is the first to change from 0 to 1 then (5.3) illustrates that the resulting signature has a coefficient of 1 corresponding to the word $(1,2)$ and a coefficient of 0 corresponding to the word $(2,1)$. If the y -coordinate is the first to change from 0 to 1 then (5.4) illustrates that the resulting signature has a coefficient of 0 corresponding to the word $(1,2)$ and a coefficient of 1 corresponding to the word $(2,1)$. Consequently the 7 real numbers forming the coefficients of the truncation of the signature to depth 2 are sufficient to determine the order. This is *independent* of the number of samples N , and dealing with 7 terms is simpler than dealing with $1 + 3N + 2N^2$ quadratic polynomials.

It is possible to construct streams by providing the raw data of Lie increments with higher order terms by specifying the width. For example, consider the stream of increments $\mathbf{x} = ((1,1,1), (1,-3,4), (0,1,2)) \in \mathcal{S}(\mathbb{R}^3)$. Consider taking width 2 and depth 2 so that the elements of a Lie element will have keys $(1,2, [1,2])$. Then the python code in Figure 6 will construct a stream from \mathbf{x} whose underlying Lie increments are width 2, depth 2.

As a final example, consider the stream $\mathbf{y} = \{(1,1,2.5,0.75,-1.5), (1.2,3.4,-2.5,0.98,0.4)\} \in \mathcal{S}(\mathbb{R}^5)$. Suppose we want to determine a Lie element of width 2 and depth 3 from the stream \mathbf{y} , which will have keys $(1,2, [1,2], [1, [1,2]], [2, [1,2]])$. Then the python code in Figure 7 will construct a stream from \mathbf{y} whose underlying Lie increments are width 2, depth 3.

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


```

>>> import numpy as np
>>> import roughpy as rp
>>> x = np.array([ [1.0,1.0,1.0] , [1.0,-3.0,4.0] , [0.0,1.0,2.0] ])
>>> context = rp.get_context(width=2,depth=2, coeffs = rp.DPReal)
>>> stream = rp.LieIncrementStream. from_increments(x, ctx = context)
>>> interval1 = rp.RealInterval(0,1)
>>> interval2 = rp.RealInterval(0,2)
>>> interval3 = rp.RealInterval(0,3)
>>> logsig1 = stream.log_signature(interval1, ctx = context)
>>> logsig2 = stream.log_signature(interval2, ctx = context)
>>> logsig3 = stream.log_signature(interval3, ctx = context)
>>> print(logsig1)
{ 1(1) 1(2) 1([1,2]) }
>>> print(logsig2)
{ 2(1) -2(2) 3([1,2]) }
>>> print(logsig3)
{ 2(1) -1(2) 6([1,2]) }

```

Figure 6: Python code to use the stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^3)$ to prescribe the depth 2 information. Includes outputs resulting from computing and printing the log signature over a range of intervals.

```

>>> import numpy as np
>>> import roughpy as rp
>>> y = np.array([ [1.0,1.0,2.5,0.75,-1.5] , [1.2,3.4,-2.5,0.98,0.4] ])
>>> context = rp.get_context(width=2,depth=3, coeffs = rp.DPReal)
>>> stream = rp.LieIncrementStream. from_increments(x, ctx = context)
>>> interval1 = rp.RealInterval(0,1)
>>> interval2 = rp.RealInterval(0,2)
>>> logsig1 = stream.log_signature(interval1, ctx = context)
>>> logsig2 = stream.log_signature(interval2, ctx = context)
>>> print(logsig1)
{ 1(1) 1(2) 2.5([1,2]) -0.75([1,[1,2]]) -1.5([2,[1,2]]) }
>>> print(logsig2)
{ 2.2(1) 4.4(2) 1.1([1,2]) -2.55667([1,[1,2]]) -7.04([2,[1,2]]) }

```

Figure 7: Python code to use the stream $\mathbf{y} \in \mathcal{S}(\mathbb{R}^5)$ to prescribe the depth 3 information. Includes outputs resulting from computing and printing the log signature over a range of intervals.

examples to explore the behaviour of path signatures and of the affect that various stream transformations (including those covered in Section 2.5) 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, if $M \in \mathbb{Z}_{\geq 1}$ then a finite collection $\{\gamma_1, \dots, \gamma_M\} \subset \mathcal{V}^1([a, b], V)$ can be summarised by the empirical measure $\delta := \frac{1}{M} \sum_{j=1}^M \delta_{\gamma_j}$. We would like to summarise distributions in an analogous manner to how the signature summarises paths. That is, loosely speaking, compress the information whilst retaining enough detail to distinguish between different distributions. For paths the signature achieves this by filtering out the parameterisation noise (cf. Section 3). Typically, summarising distributions requires restricting the class of distribution considered by imposing additional constraints (cf. Brownian Motion, Markov processes etc). It turns out that the signature allows one to summarise *all* distributions without restriction.

Let $p \in \mathbb{R}_{\geq 1}$ and let $\mathcal{P}\Omega_p(V)$ denote the space of (Borel) probability measures on the space of p -rough paths $\Omega_p(V)$. Then the *Expected Signature* is the map $\mathbb{S}_p : \mathcal{P}\Omega_p(V) \rightarrow T((V))$ defined by

$$\mathbb{S}_p(\mu) := \mathbb{E}_\mu[S(X)] = \prod_{n=0}^{\infty} \mathbb{E}_\mu[S^n(X)]. \quad (6.1)$$

The expected signature defined in (6.1) allows us to define the expected signature of a stochastic process.

For this purpose, suppose that X_t is a stochastic process on V for $t \in [0, 1]$ under a probability space $(\mathcal{N}, \mathbb{P})$. First suppose that $p \in [1, 2)$ and that for \mathbb{P} -almost every $\omega \in \mathcal{N}$ the path $[0, 1] \rightarrow V$ defined by $s \mapsto X_s(\omega)$ is in $\mathcal{V}^p([0, 1], V)$. Let $Q : \mathcal{N} \rightarrow \mathcal{V}^p([0, 1], V)$ be the map taking $\omega \in \mathcal{N}$ to the path $(s \mapsto X_s(\omega))_{s \in [0, 1]}$. The push-forward $Q_*\mathbb{P}$ of the measure \mathbb{P} is an element in $\mathcal{P}\mathcal{V}^p([0, 1], V)$. Recalling from Subsection 2.9 that $p \in [1, 2)$ means that p -rough paths in $\Omega_p(V)$ are the signatures of paths in $\mathcal{V}^p([0, 1], V)$, we may view $Q_*\mathbb{P}$ as an element in $\mathcal{P}\Omega_p(V)$. Thus we can make sense of the *expected signature of X_t* as $\mathbb{S}_p(Q_*\mathbb{P})$, with it being common to abuse notation and write $\mathbb{S}_p(X)$. When the value of $p \in [1, 2)$ is either clear from the context or not of particular importance, it is common to drop it from the notation and write $\mathbb{S}(X)$ for the expected signature of the process X_t .

If the sample paths of the stochastic process are only of finite p -variation for $p \geq 2$ then making sense of the expected signature of the process is more involved. Suppose that this is now the case, and assume that for \mathbb{P} -almost every $\omega \in \mathcal{N}$ there are, for $i \in \{2, \dots, \lfloor p \rfloor\}$, functions $A[i](\omega) : \Delta_{[a, b]} \rightarrow V^{\otimes i}$ such that the functional $\mathbf{X}(\omega) : \Delta_{[a, b]} \rightarrow T^{(\lfloor p \rfloor)}(V)$ defined for $(s, t) \in \Delta_{[a, b]}$ by

$$\mathbf{X}_{s, t}(\omega) := \left(1, X_t(\omega) - X_s(\omega), A_{s, t}[2](\omega), \dots, A_{s, t}[\lfloor p \rfloor](\omega)\right) \in T^{(\lfloor p \rfloor)}(V) \quad (6.2)$$

satisfies that $\mathbf{X}(\omega) \in \Omega_p(V)$. Let $Q : \mathcal{N} \rightarrow \Omega_p(V)$ be the map taking $\omega \in \mathcal{N}$ to the p -rough path $\mathbf{X}(\omega) \in \Omega_p(V)$. The push-forward $Q_*\mathbb{P}$ of the measure \mathbb{P} is an element in $\mathcal{P}\Omega_p(V)$. Consequently we can make sense of the *expected signature of X_t* as $\mathbb{S}_p(Q_*\mathbb{P})$ with it being common to abuse notation and write $\mathbb{S}_p(X)$. When the value of $p \geq 2$ is either clear from the context or not of particular importance, it is common to drop it from the notation and write $\mathbb{S}(X)$ for the expected signature of the process X_t .

In this latter case we observe that the increments of the process X_t are no longer sufficient to determine its expected signature. Instead, we are required to additionally prescribe its higher order properties such as its area. A standard Brownian motion $B : [0, T] \rightarrow \mathbb{R}^d$ for some $T > 0$ provides an instructive example. Given a $p \in (2, 3)$ it is known that the process $t \mapsto B_t$ possesses finite p -variation. Consequently, in order to define its expected signature we must first make a choice of depth 2 term; that is, we need to prescribe a suitable function $\Delta_{[0, T]} \rightarrow T^{(2)}(\mathbb{R}^d)$.

Both the Itô 2-multiplicative functional (cf. (2.70)) and the canonical Brownian rough path (cf. (2.75)) defined in Subsection 2.9 provide choices of depth 2 terms that result in distinct p -rough paths whose depth 1 terms are given by the increments $B_t - B_s$ of the Brownian motion. The resulting expected signatures are both candidates for the expected signature of the Brownian motion. Without making a choice of depth 2 term there is no canonical way to choose which one should be taken to be the expected signature of the Brownian motion B_t .

The expected signature of a process X_t is the natural generalisation of the moments of the process X_t . Moreover, in [Bonnier and Oberhauser, 2020], 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 [Fawcett, 2003], 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 $S(X)$ [Chevyrev, 2014]. 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 [Chevyrev and Lyons, 2016, Chevyrev and Oberhauser, 2022].

For Brownian motion B_t with Lévy area (i.e. the canonical Brownian rough path defined in Subsection 2.9; cf. (2.75)) 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 [Lyons and Ni, 2015], where the authors established that the radius of convergence was strictly positive. In [Chevyrev and Lyons, 2016] 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 [Boedihardjo et al., 2021] 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 $T^{(n)}(V)$ [Lyons and Victoir, 2004, Litterer and Lyons, 2007, Lyons and Yang, 2013]. Unlike the *Martingale Problem* approach originating in [Stroock and Varadhan, 1997], there is no restriction on the probability measures considered.

A result from [Chevyrev and Oberhauser, 2022] 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 [Chevyrev and Oberhauser, 2022] that when restricted to compact subsets, the expected signature is injective up-to tree-like equivalence (see Appendix B of [Chevyrev and Oberhauser, 2022] 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 $\mathcal{PV}^1([a, b], V)$. Let $\mathcal{C}([a, b], V) \subset \mathcal{V}^1([a, b], V)$ denote the subset resulting from removing all tree-like paths.

The work [Lemercier et al., 2021] establishes a universality property for the expected signature analogous to the universality property of the signature itself (cf. Theorem 3.3). 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 $C^0(\mathcal{PK})$, where $\mathcal{K} \subset \mathcal{C}([a, b], V)$ is compact and \mathcal{PK} denotes the space of Borel probability measures on \mathcal{K} , 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 $T((V))$.

Before introducing the lift of the expected signature map to the so-called *Pathwise Expected Signature* $\Phi_{\text{Path}} : \mathcal{PC}([a, b], V) \rightarrow \mathcal{V}^1([a, b], T((V)))$, we first fix a choice of topology on the space of Borel probability measures \mathcal{PK} on a compact subset $\mathcal{K} \subset \mathcal{C}([a, b], V)$. For this purpose we consider the following notion of weak convergence in \mathcal{PK} .

Let $C^0(\mathcal{K})$ denote the space of real-valued continuous functions $\mathcal{K} \rightarrow \mathbb{R}$. Then any $\nu \in \mathcal{PK}$ can be viewed as a bounded linear functional $C^0(\mathcal{K}) \rightarrow \mathbb{R}$ by defining $\nu[\varphi] := \int_{\mathcal{K}} \varphi(z) d\nu(z)$ for $\varphi \in C^0(\mathcal{K})$. Then a sequence $\{\mu_n\}_{n=1}^{\infty} \subset \mathcal{PK}$ is said to weakly converge to $\mu \in \mathcal{PK}$ if for every $f \in C^0(\mathcal{K})$ we have that $\mu_n[f] \rightarrow \mu[f]$ as $n \rightarrow \infty$. It is established in [Lemercier et al., 2021] (cf. Theorem A.1) that the expected signature determines a weakly continuous map $\mathcal{PK} \rightarrow T((V))$ where $T((V))$ is equipped with the topology induced by $\langle \cdot, \cdot \rangle_{T((V))}$ introduced in Subsection 2.7.

The *Pathwise Expected Signature* is the map $\Phi_{\text{Path}} : \mathcal{PC}([a, b], V) \rightarrow \mathcal{V}^1([a, b], T((V)))$ defined by

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

It is established in [Lemercier et al., 2021] (cf. Theorem A.4) that if $\mathcal{K} \subset \mathcal{C}([a, b], V)$ is compact, then the pathwise expected signature determines a weakly continuous map $\mathcal{PK} \rightarrow \mathcal{V}^1([a, b], T((V)))$.

The coordinate iterated integrals forming the components of $S_{a,b}(\Phi_{\text{Path}}(\mu))$ then span $C^0(\mathcal{PK})$ for compact subsets $\mathcal{K} \subset \mathcal{C}([a, b], V)$. That is, any weakly continuous function on \mathcal{PK} can be uniformly well-approximated by a linear combination of the terms in the signature of the pathwise expected signature.

Theorem 6.1 (Expected Signature Universality; Theorem 3.2 in [Lemercier et al., 2021]). *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 weakly continuous function $F : \mathcal{PK} \rightarrow \mathbb{R}$. Then for any $\varepsilon > 0$ there exists a truncation level $m \in \mathbb{N}$ and real coefficients $\{\alpha_k(\mathbf{J}) \mid k \in \{0, \dots, m\} \text{ and } \mathbf{J} \in \{1, \dots, d\}^k\} \subset \mathbb{R}$ such that for any $\mu \in \mathcal{PK}$ 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 3.2). 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 . 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

Path signatures live 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 3.2) 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 $S^{(N)}(x) := \Pi_N(S(x))$.

From a theoretical point of view this is not an issue. Indeed, on compact subsets, we know that provided a sufficiently large truncation depth N is chosen, we are able to approximate continuous functions by linear combinations of the coordinate iterated integrals of the truncated to depth N signatures (cf. Theorem 3.3). However, from a practical perspective there is a potential problem.

The theoretical guarantees provide no upper bound on the truncation depth required. In particular, a large value of N may be required to adequately capture continuous functions with non-trivial dependence on higher order signature terms. This can become problematic due to the exponential growth in the number of components in each term of a signature with respect to the depth. For example, suppose V is d -dimensional for some integer $d \in \mathbb{Z}_{\geq 2}$. Then given an integer $k \in \mathbb{Z}_{\geq 1}$, the depth k term of a signature of a path in V consists of d^k components. Consequently, the truncation of the signature to depth N involves

$$1 + d + \dots + d^N = \frac{d^{N+1} - 1}{d - 1}$$

which quickly becomes intractable for moderately sized values of d and N . Despite signatures compressing the information of a path down to a countable graded collection of statistics, further compression is required to

enable directly learning a suitable linear combination.

Introducing a suitable augmentation of the dataset *before* taking the signature offers a solution to this problem. Recall (cf. Section 2.5) that by augmenting the dataset $\Omega_{\mathcal{S}(V)}$, we mean considering a map $\Theta : V \rightarrow W$ for some Banach space W . Any stream $\mathbf{x} \in \mathcal{S}(V)$ can be transformed to be a stream in $\mathcal{S}(W)$ via the element-wise application of the map Θ . Hence we can work with the augmented dataset $\Omega_{\mathcal{S}(W)} \subset \mathcal{S}(W)$ resulting from using Θ to transform each stream in $\Omega_{\mathcal{S}(V)}$ to a stream in $\mathcal{S}(W)$.

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 signature of the stream $\Theta(\Omega_V)$ truncated to a manageable depth. A common situation is that the Banach space V is \mathbb{R}^d for some $d \in \mathbb{Z}_{\geq 1}$, and that the Banach space U is chosen to be \mathbb{R}^e for some $e \in \mathbb{Z}_{\geq 1}$ so that $W = \mathbb{R}^{d+e}$.

We have already seen several explicit examples of possible augmentations in Section 2.5, and each augmentation offers different benefits. A Lead-Lag augmentation aides the study of a streams variance via its signature (see [Chevyrev and Kormilitzin, 2016]), the invisibility-reset augmentation (cf. (2.21)) 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. (2.17)) 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 [Chevyrev and Oberhauser, 2022, Király and Oberhauser, 2019] for example. It is proposed in [Kidger et al., 2019] 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 [Géron, 2022], [Heaton, 2011], [Gurney, 1997], 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 3.3). 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 [Kidger et al., 2019] 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 [Kidger et al., 2019].

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 [Lyons et al., 2007]. Two simple architectures involving signature layers are shown in Figure 8, and full details of the proposed models can be found in [Kidger et al., 2019].




Figure 8: Two simple architectures illustrating the possible inclusion of a signature layer - Figure 1 in [Kidger et al., 2019]. Reproduced with permission. Copyright held by authors of [Kidger et al., 2019].

The required computation and back-propagation of the signature transform is efficiently managed by the *Signatory* package [Kidger and Lyons, 2020]. The resulting model obtained in [Kidger et al., 2019] is compared with 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

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

for some realisation of B^H . Empirical evidence suggests fractional Brownian motions are closely linked with financial market data [Gatheral et al., 2018], and estimating the Hurst parameter is a non-trivial task [Lacasa et al., 2009].

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 $[0.2, 0.8]$. 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 Table 1.

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 1: Final test MSE for the different models - Variant of Table 1 in [Kidger et al., 2019]

The *Rescaled Range* method is a mathematically derived method involving no learning, see [Hurst, 1951]. 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 out perform 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 [Kidger et al., 2019].

8. Signature Kernel

The signature provides a universal feature map embedding the dataset $\Omega_{\mathcal{S}(V)}$ into the tensor algebra $T((V))$. By considering the inner product we equipped the tensor algebra $T((V))$ with in Section 2.7, we can 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 $\mathcal{H}(V) := \overline{T((V))}$ is a Hilbert space, meaning that the signature provides us with a kernel. Recall that we consider a labelled data set $\Omega \subset \mathcal{S}(V) \times \mathbb{R}$ given, for some $M \in \mathbb{Z}_{\geq 1}$, by $\Omega = \{(\mathbf{x}_k, y_k) \in \mathcal{S}(V) \times \mathbb{R} \mid k \in \{1, \dots, M\}\}$, and we take $\Omega_{\mathcal{S}(V)} \subset \mathcal{S}(V)$ to be the projection of Ω onto $\mathcal{S}(V)$. Moreover we assume that each stream $\mathbf{x} \in \Omega_{\mathcal{S}(V)}$ consists of incremental data. For each $k \in \{1, \dots, M\}$ let Γ_k denote the contour resulting from the concatenation of the entries in the stream \mathbf{x}_k , and choose a parameterisation $\gamma_k \in \mathcal{V}^1([0, 1], V)$ of Γ_k .

Then we can consider the function $K_{\Omega_{\mathcal{S}(V)}} : \Omega_{\mathcal{S}(V)} \times \Omega_{\mathcal{S}(V)} \rightarrow \mathbb{R}_{\geq 0}$ given by

$$K_{\Omega_{\mathcal{S}(V)}}(\mathbf{x}_i, \mathbf{x}_j) := K_{\gamma_i, \gamma_j}(1, 1) = \langle S_{0,1}(\gamma_i), S_{0,1}(\gamma_j) \rangle_{T((V))} \quad (8.2)$$

for $i, j \in \{1, \dots, M\}$. In [Salvi et al., 2021] the authors consider the use of the kernel function $K_{\Omega_{\mathcal{S}(V)}}$, 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* [Lees, 1960]), the authors introduce a method to use the full signature transform for regression tasks *without* any truncation via use of the kernel function $K_{\Omega_{\mathcal{S}(V)}}$. The primary aim for the remainder of this section is to summarise this approach developed in [Salvi et al., 2021].

Kernels on truncated signatures are considered in [Király and Oberhauser, 2019] via the development of a technique for transforming a kernel on V to a kernel on $\mathcal{S}(V)$. The approach followed in [Király and Oberhauser, 2019] is to use truncated kernels (in which the full signature is replaced by its truncation to some finite depth) 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 [Toth and Oberhauser, 2020].

In [Salvi et al., 2021] 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 [Salvi et al., 2021]). *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*

$$K_{x,y}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{T((V))}. \quad (8.3)$$

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

$$\frac{\partial^2 K_{x,y}}{\partial s \partial t} = \langle \dot{x}_s, \dot{y}_t \rangle K_{x,y} \quad (8.4)$$

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 (8.4) is an example of a *Goursat problem*, for which existence and uniqueness results originate in [Lees, 1960]. Consequently, for sufficiently regular paths, the existence and uniqueness of K solving (8.3) is guaranteed (cf. Theorem 3.1 in [Salvi et al., 2021]). The content of Theorem 8.1 makes sense for rougher paths. That is, the PDE (8.4) can be made sense of for paths x and y of lower regularity by considering its integral form

$$K_{x,y}(s, t) = 1 + \int_a^s \int_c^t \langle S_{a,p}(x), S_{c,q}(y) \rangle_{T((V))} \langle dx_p, dy_q \rangle_V \quad (8.5)$$

and making an appropriate definition for the quantity $\langle dx_p, dy_q \rangle_V$. This is done in section 4 in [Salvi et al., 2021] for the setting that the paths x and y are geometric rough paths (cf. Definition 2.11 in this article). The main technicality is using the theory of integrating one-forms along rough paths to make sense of the double integral in (8.5) in this rougher setting. A review of the theory of integrating one-forms along rough paths can be found in Appendix B of [Salvi et al., 2021]; a more thorough presentation of this theory can be found in section 4 of [Lyons et al., 2007]. The full details of how to make sense of the *rough double integral* can be found in section 4 of [Salvi et al., 2021].

Returning to the setting of paths of bounded variation, numerical schemes to approximate $K_{x,y}$ are developed in [Salvi et al., 2021], 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 (8.4) becomes

$$\frac{\partial^2 K_{x,y}}{\partial s \partial t} = C K_{x,y} \quad (8.6)$$

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, (8.6) reads

$$K_{x,y}(s, t) = K_{x,y}(s, v) + K_{x,y}(u, t) - K_{x,y}(u, v) + C \int_u^s \int_v^t K_{x,y}(r, w) dr dw \quad (8.7)$$

for $(s, t), (u, v) \in \mathcal{D}_{ij}$ with $u \leq s$ and $v \leq t$. Approximations of the double-integral in (8.7) can be used to develop finite difference schemes to approximate $K_{x,y}$. The relatively simple approximation

$$\int_u^s \int_v^t K_{x,y}(r, w) dr dw \approx \frac{1}{2} (K_{x,y}(s, v) + K_{x,y}(u, t)) (s - u)(t - v), \quad (8.8)$$

is used in [Salvi et al., 2021] 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 (8.8) can be used to define numerical schemes for (8.4) 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

$$P_\lambda \cap ([u_i, u_{i+1}] \times [v_j, v_{j+1}]) = \bigcup_{k,l=0}^{2^\lambda} \left\{ \left(u_i + \frac{k}{2^\lambda} (u_{i+1} - u_i), v_j + \frac{l}{2^\lambda} (v_{j+1} - v_j) \right) \right\}. \quad (8.9)$$

In [Salvi et al., 2021] the authors define the finite difference scheme

$$\begin{aligned} \hat{K}(s_{i+1}, t_{j+1}) &= \hat{K}(s_{i+1}, t_j) + \hat{K}(s_i, t_{j+1}) - \hat{K}(s_i, t_j) + \\ &\quad \frac{1}{2} \langle x_{s_{i+1}} - x_{s_i}, y_{t_{j+1}} - y_{t_j} \rangle (\hat{K}(s_{i+1}, t_j) + \hat{K}(s_i, t_{j+1})), \end{aligned} \quad (8.10)$$

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

$$P_\lambda = \bigcup_{i=0}^{m-1} \bigcup_{j=0}^{n-1} \bigcup_{k,l=0}^{2^\lambda-1} \left\{ \left(u_i + \frac{k}{2^\lambda} (u_{i+1} - u_i), v_j + \frac{l}{2^\lambda} (v_{j+1} - v_j) \right) \right\} =: \bigcup_{i=0}^{2^\lambda n} \bigcup_{j=0}^{2^\lambda m} \{(s_i, t_j)\}. \quad (8.11)$$

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 [Salvi et al., 2021]). *Assume that x and y are piecewise linear paths as above satisfying that*

$$\sup_{I \times J} |\langle \dot{x}_s, \dot{y}_t \rangle| \leq M \quad (8.12)$$

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 (8.10) on P_λ satisfies

$$\sup_{I \times J} |K_{x,y}(s, t) - \hat{K}(s, t)| \leq \frac{C_1}{2^{2\lambda}}. \quad (8.13)$$

In [Salvi et al., 2021] 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 9.




Figure 9: Example of error distribution of $k_{x,y}(s, t)$ on the grids P_0 , P_1 and P_2 —Figure 2 in [Salvi et al., 2021]. Copyright ©2021 Society for Industrial and Applied Mathematics. Reprinted with permission. All rights reserved.

It is noted in [Salvi et al., 2021] that the “PDE structure” allows for efficient GPU implementation of the scheme via parallelization. Figure 10, which is Figure 3 in [Salvi et al., 2021], illustrates the complexity reduction from quadratic on CPU to linear on GPU.




Figure 10: 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. Reproduction of Figure 3 in [Salvi et al., 2021]. Copyright ©2021 Society for Industrial and Applied Mathematics. Reprinted with permission. All rights reserved.

In [Salvi et al., 2021], the untruncated signature kernel is compared against other kernels via the task of classification with Support Vector Machine (SVM) classifiers [Vapnik, 1998]. SVM classifiers have been

successfully used in medicine [Furey et al., 2000], image retrieval [Tong and Chang, 2001] and text classification [Tong and Koller, 2001], 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 $\{(x_i, y_i) \mid i \in \{1, \dots, n\}\}$. For binary classification we have $y_i \in \{-1, 1\}$ and the binary SVM classification algorithm aims at solving the minimisation problem

$$\min_{f \in H_k} \sum_{i=1}^n L(y_i, f(x_i)) + \lambda \|f\|_{H_k} \quad (8.14)$$

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

$$f_*(x) := \text{sign} \left(\alpha_0 + \sum_{i=1}^n \alpha_i y_i k(x, x_i) \right) \quad (8.15)$$

where $\alpha_0, \dots, \alpha_n$ are scalar coefficients solving a certain quadratic programming problem, see [Scholkopf and Smola, 2018] for full details.

Selecting an appropriate kernel when X consists of multivariate time-series is difficult [Sapankevych and Sankar, 2009]. 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 SVM classifier equipped with a range of different kernels on various multivariate time-series UEA datasets [Bagnall et al., 2017] are presented in Table 2. The *global alignment kernel* (GAK) [Cuturi, 2011] 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 [Salvi et al., 2021].

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 2: Test set classification accuracy (in %) on UEA multivariate time series datasets - Table 1 in [Salvi et al., 2021]

Subsequent to the appearance of [Salvi et al., 2021], numerous works have considered the development and application of signature kernel techniques within a wide range of settings. We end this section with a far-from-exhaustive summary of some of the further developments of signature kernel techniques.

General Signature Kernels (GSKs) have been introduced in [Cass et al., 2021]. Recall that for $A, B \in T((V))$ we have that $\langle A, B \rangle_{T((V))} = \sum_{n=0}^{\infty} \langle \pi_n(A), \pi_n(B) \rangle_{V^{\otimes n}}$ (cf. (2.37)). Loosely, the GSKs proposed in [Cass et al., 2021] are obtained by choosing coefficients $\{\phi(n)\}_{n=0}^{\infty} \subset \mathbb{C}$ that are *not* required to be 1, and subsequently replacing the inner product $\langle \cdot, \cdot \rangle_{T((V))}$ by $\langle \cdot, \cdot \rangle_{\phi}$ which is defined for $A, B \in T((V))$ by $\langle A, B \rangle_{\phi} := \sum_{n=0}^{\infty} \phi(n) \langle \pi_n(A), \pi_n(B) \rangle_{V^{\otimes n}}$. 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

$$K_{x,y}^{\phi}(s, t) := \langle S_{a,s}(x), S_{c,t}(y) \rangle_{\phi} = \sum_{n=0}^{\infty} \phi(n) \langle S_{a,s}^n(x), S_{c,t}^n(y) \rangle_{V^{\otimes n}} \quad (8.16)$$

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 [Cass et al., 2021]. The extension of this analysis allows the authors of [Cass et al., 2021] 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.

In a different direction the signature kernel techniques developed in [Salvi et al., 2021] are central to the work on Optimal Stopping Problems in [Horvath et al., 2023]. A major achievement of [Horvath et al., 2023] is the development of a rigorous mathematical framework in which an Optimal Stopping Problem can be recast as a higher order kernel mean embedding regression based on the notions of higher rank signatures of measure-valued paths and adapted topologies. This overcomes the issue in Optimal Stopping Problems arising from mathematical finance that the corresponding value functions are, in general, discontinuous with respect to the weak topology. Full details may be found in [Horvath et al., 2023].

The final use of the signature kernel techniques developed in [Salvi et al., 2021] that we mention here is in the development of an efficient method for the computation of a sparse collection of iterated integrals [Shmelev and Salvi, 2024]. The authors of [Shmelev and Salvi, 2024] introduce a technique to effectively isolate specific groups of signature coefficients including, in particular, a singular coefficient at any depth. The empirical effectiveness of this approach is demonstrated via the task of the construction of N -step Euler schemes for sparse CDEs [Shmelev and Salvi, 2024].

9. Neural CDEs and Neural RDEs

In this section we discuss *Neural Controlled Differential Equations* (Neural CDEs) [Kidger et al., 2020] and their subsequent development to *Neural Rough Differential Equations* (Neural RDEs) [Morrill et al., 2021b]. 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 [Chen et al., 2018]. 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 [Kidger et al., 2020].

In [Kidger et al., 2020] the authors resolve this issue by using 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 $\frac{dX_s}{ds} ds$. In Appendix A of [Kidger et al., 2020] it is justified 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 between Neural CDEs and Neural ODEs is that the equation (9.2) for Neural CDEs is driven by the data process X , whilst the corresponding 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 11, which is Figure 1 in [Kidger et al., 2020].




Figure 11: **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 [Kidger et al., 2020]. Reproduced with permission. Copyright held by authors of [Kidger et al., 2020].

In Appendix B in [Kidger et al., 2020] the authors establish that Neural CDEs are universal approximators. The precise statement may be found as Theorem B.14 in [Kidger et al., 2020]; 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 3.3 in this article) and that the path signature solves a CDE (cf. (4.2)). Full details may be found in Appendix B of [Kidger et al., 2020].

The efficacy of Neural CDEs is demonstrated in [Kidger et al., 2020] via the task of written character classification. The CharacterTrajectories dataset from the UEA time series classification archive [Bagnall et al., 2018] 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 [Kidger et al., 2020] 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, completely observed time series. The Neural CDE model is compared with several ODE and RNN based models, details of which may be found in [Kidger et al., 2020]. The results are summarised in Table 3 below, which is a variant of Table 1 in [Kidger et al., 2020]. 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.

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 3: Test accuracy (mean% \pm std%, computed across five runs) on CharacterTrajectories - Table 1 in [Kidger et al., 2020]

The performance of Neural CDEs is investigated through further numerical experiments in [Kidger et al., 2020]. 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 [Kidger et al., 2020]. It is also explained in [Kidger et al., 2020] 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 [Kidger et al., 2020]. The subsequent work [Morrill et al., 2021a] 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 [Morrill et al., 2021a] for full details.

The Neural CDE approach developed in [Kidger et al., 2020] 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 [Morrill et al., 2021b]. In [Morrill et al., 2021b] 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 [Morrill et al., 2021b]. 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 hyper-parameters of the model. For a chosen depth hyper-parameter $N \in \mathbb{Z}_{\geq 1}$ 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. (4.8)) 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 12 below gives a high level comparison of the CDE and RDE based approaches; it appears as Figure 1 in [Morrill et al., 2021b]



(a) CDE approach proposed in [Kidger et al., 2020] in which data is smoothly interpolated and pointwise derivative information is used to drive the CDE.

(b) RDE approach proposed in [Morrill et al., 2021b] 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 12: Figure 1 in [Morrill et al., 2021b]. Reproduced with permission. Copyright held by authors of [Morrill et al., 2021b].

Numerical experiments are conducted in [Morrill et al., 2021b] 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 [Bagnall et al., 2017]. 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 $\{2^j \mid j \in \{1, \dots, 10\}\}$. A Neural CDE model and an ODE-RNN model (introduced in [Rubanova et al., 2019]) are used for baseline comparison. Precise details of hyper-parameter selection, optimisers, and other model parameter choices may be found in Appendix C of [Morrill et al., 2021b].

The experiment is run three times for each model and each hyper-parameter combination, with the mean and standard deviation of the accuracy recorded. The results are presented in Table 4 below, which is a version of Table 1 in [Morrill et al., 2021b]. For brevity, the results are only displayed for a sub-selection of step sizes considered; the full results can be found in Table 7 in Appendix D of [Morrill et al., 2021b]. Amongst the results included in Table 4 we observe that the best accuracy results are returned by Neural RDE models, and that the accuracy achieved is a noticeable improvement over that resulting from either of ODE-RNN or Neural CDE models. This improvement over Neural CDE performance comes at the cost of Neural RDE models being more memory intensive. However, Neural RDE models remain significantly less memory intensive than ODE-RNN models.

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 4: Test accuracy (mean% \pm std%, computed across three repeats) and memory usage on EigenWorms - Part of Table 1 in [Morrill et al., 2021b]

In [Kidger et al., 2021a] Neural CDEs are combined with *Neural Stochastic Differential Equations* (Neural SDEs) (see [Tzen and Raginsky, 2019, Li et al., 2020, Hodgkinson et al., 2020], 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 pre-specified 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 [Kidger et al., 2021a]. Moreover, the follow-up work [Kidger et al., 2021b] introduces several technical innovations to improve both model performance and training speed for Neural SDEs.

Neural RDEs reduce the computational cost of the log-ODE method by learning the vector field \tilde{f} (cf. (4.5) in Section 4) as a neural network rather than computing it directly. The paper [Walker et al., 2024] introduces *Log-Neural Controlled Differential Equations* (Log-NCDEs) which are Neural RDEs that do *not* ignore the structure of the vector field \tilde{f} . Indeed Log-NCDEs directly compute the vector field \tilde{f} via the iterated Lie brackets of the original vector field f . Whilst Log-NCDEs are necessarily more costly than Neural CDEs/RDEs, it is shown in [Walker et al., 2024] that this exploitation of the underlying regularity/structure of the vector fields is demonstrated to achieve a higher average test set accuracy than Neural CDEs, Neural RDEs, S5, and the linear recurrent unit on a range of multivariate time series classification benchmarks; the interested reader is directed to [Walker et al., 2024] for full precise details.

The **Neural CDE** notebook contains an introduction to the technology developed in both [Kidger et al.,

2020] and [Morrill et al., 2021b]. 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 [Aldeneh and Provost, 2017].

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) [Neumann and Vu, 2017, Mirsamadi et al., 2017, Tao and Liu, 2018, Han et al., 2018]. 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 [Wang et al., 2019], 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 [Mirsamadi et al., 2017].

Inspired by [Yang et al., 2016], 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.

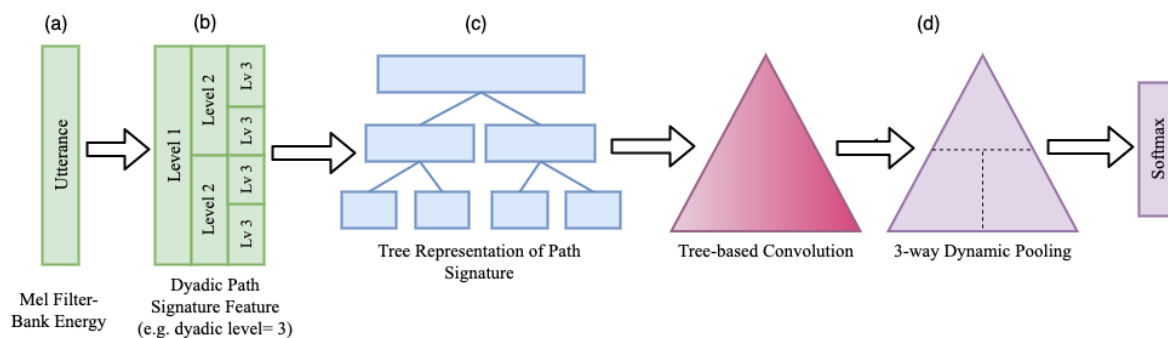


Figure 13: Overview of the proposed SER system - Figure 1 in [Wang et al., 2019]. Reprinted with permission.

Figure 13 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 [Wang et al., 2019] 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 [Busso et al., 2008]. 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 [Wang et al., 2019], 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 [Eyben et al., 2010] 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 hyper-parameters used are listed in table 1 of [Wang et al., 2019]. 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 5.

Model	UA	WA
COVAREP [Ghosh et al., 2016]	51.84	49.64
DNN-ELM [Lee and Tashev, 2015]	52.13	57.91
LSTM (Speech) [Ghosh et al., 2016]	51.85	51.94
LSTM (Glottal) [Ghosh et al., 2016]	54.56	52.82
Attentive CNN [Neumann and Vu, 2017]	56.83	61.95
PTS-CNN [Wang et al., 2019]	53.03	58.90

Table 5: Model Performances - Table 2 in [Wang et al., 2019]

The relatively simple Path-Tree-Signature based CNN (PTS-CNN) model in [Wang et al., 2019] 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 [Eyben et al., 2010] 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 [Eyben et al., 2015]. 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 [Ghosh et al., 2016]).

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 [Perez Arribas et al., 2018]. 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 [Perez Arribas et al., 2018], signature-based machine learning models are proposed to analyse data obtained from a clinical study in [Tsanas et al., 2016]. 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 14 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 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 score 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.

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 15. 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 [Tsanas et al., 2016] 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 [Perez Arribas et al., 2018].

For the mood prediction, mean absolute error (MAE) regression was used on the input-output pairs




Figure 14: 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 path — Figure 1 in [Perez Arribas et al., 2018]. Unaltered. Licensed under a [Creative Commons Attribution 4.0 International License](#).

$\{(S^{(n)}(R_i), Y_i)\}$. 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 Table 6. More detailed performance results for the signature-based method can be found in Table 3 in [Perez Arribas et al., 2018]

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

Table 6: Comparison of mood prediction models — Summary of Table 3 in [Perez Arribas et al., 2018]

The subsequent work [Wang et al., 2020] 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 [Perez Arribas et al., 2018, Tsanas et al., 2016] 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




Figure 15: 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 [Perez Arribas et al., 2018]. Unaltered. Licensed under a [Creative Commons Attribution 4.0 International License](#).

- 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 [Wang et al., 2020] 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 [Wang et al., 2020]. 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 [Wang et al., 2020].

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 [Loranger et al., 1994]) 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.

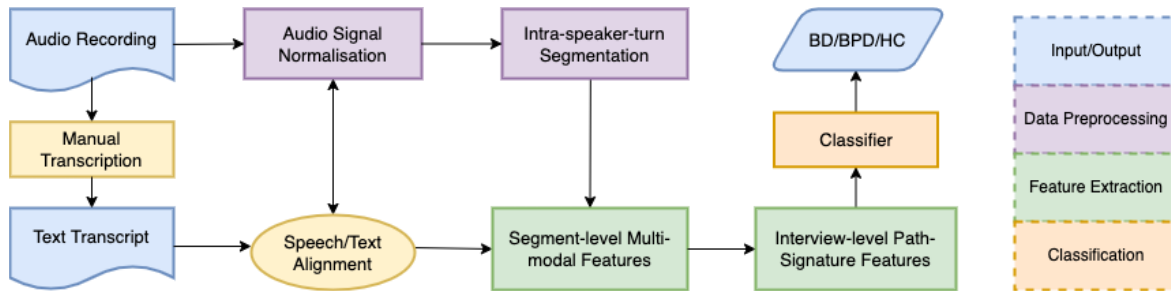


Figure 16: Data pipeline including three stages of data preprocessing (in purple), feature extraction (in green) and classification (in orange). Variant of Figure 2 in [Wang et al., 2020]. Reprinted with permission.

The results for the classification task are summarised in Table 7. 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.

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 7: Average AUROC for each binary classification — Table 3 in [Wang et al., 2020]

Ablation experiments were conducted to examine the affect of removing each feature type. The results are summarised in Table 8. 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 [Wang et al., 2021] investigates the use of the invisibility reset transform (cf. (2.21)) to effectively model the conversation dynamics by enhancing the signature representation of conversation speech. Further, the subsequent work [Wu et al., 2022] uses an approach based upon the log signature to overcome the common issue of missing responses. The central idea is to record the omission of data as a new channel, and subsequently summarise the resulting streams via the log signature.

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.

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 8: Feature ablation results (AUROC) for each task; p -value used for feature selection: * * < 0.001 and * < 0.005. Table 4 in [Wang et al., 2020].

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 [Schmitter et al., 2015, Sørensen et al., 2017, Westman et al., 2011] for example. In [Moore et al., 2019], 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 [Moore et al., 2019]. 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 [Moore et al., 2019] 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 [Moore et al., 2019] 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 [Moore et al., 2019]; *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 [Moore et al., 2019] 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 [Moore et al., 2019] for full details.

The features selected for each classification task are summarised in Table 9 which is a version of Table 2 in [Moore et al., 2019].

The notation used in Table 9 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 [Moore et al., 2019].

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 9: Comparison of feature sets selected for each classification task in [Moore et al., 2019]—Table 2 in [Moore et al., 2019]

The resulting classifiers achieve at least 90% accuracy on both classification tasks considered on the test data, see Table 3 in [Moore et al., 2019] 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 [Morrill et al., 2020]. In America alone it is thought to be responsible for one in three hospital deaths [Morrill et al., 2020], and the cost of admission and patient care has been estimated to exceed \$41.5 billion [Buchman et al., 2020a, Buchman et al., 2020b, Buchman et al., 2020c]. 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 [Kumar et al., 2006]. Early detection of sepsis is evidently crucial to improving sepsis management and mortality rates.

The works [Morrill et al., 2019, Morrill et al., 2020] 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 [Singer et al., 2016] is used to denote the onset of sepsis. The development of the model in [Morrill et al., 2019] 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 [Reyna et al., 2020]. The algorithm proposed in [Morrill et al., 2019] was the first placed entry in the official phase of the challenge [Morrill et al., 2020].

The aim of this section is to present an overview of the method developed in [Morrill et al., 2019]. In addition to [Morrill et al., 2019], we will use the follow up work [Morrill et al., 2020]. The article [Morrill et al., 2020] is an extension of the Computing in Cardiology conference proceedings paper [Morrill et al., 2019] 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 [Reyna et al., 2020]) 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 [Singer et al., 2016].

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 [Morrill et al., 2020]. The following hand-crafted features derived from the data are additionally added (see Table 1 in [Morrill et al., 2019]).

- **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 [Singer et al., 2016]. 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 hyper-parameters to be optimised during training.

The challenge data was pre-labelled with the value 1 at any location of sepsis occurrence or predefined window around sepsis onset and zero otherwise. The method proposed in [Morrill et al., 2019] 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 hyper-parameter optimisation. Precise details may be found in [Morrill et al., 2020]. The final values of the parameters can be found in Table 2 in [Morrill et al., 2019]. 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 hyper-parameter 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 [Morrill et al., 2020]. The utility scores for the model using the final hyper-parameters may be found in Table 1 in [Morrill et al., 2020].

The impact of the various features on the models performance are explored in [Morrill et al., 2020]. 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 [Morrill et al., 2020]. 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 [Morrill et al., 2019] 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 [Morrill et al., 2020]. 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 [Morrill et al., 2020].

Returning to the training set of 40 336 patients, in [Morrill et al., 2020] 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 confusion matrix presented in Table 10.

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

Table 10: 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 originates in Figure 2 in [Morrill et al., 2020].

The signature-based model presented in [Morrill et al., 2019, Morrill et al., 2020] 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 [Morrill et al., 2019, Morrill et al., 2020]. 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 work through a version of the methodology of [Morrill et al., 2019, Morrill et al., 2020] 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 [Peters et al., 2018] and BERT [Kenton and Toutanova, 2019] 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 [Alsentzer et al., 2019, Peng et al., 2019, Si et al., 2019, Huang et al., 2019]. 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 [Gu et al., 2021].

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 [Biyong et al., 2020], 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 [Biyong et al., 2020] 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 [Biyong et al., 2020].

The main extension to the Transformer architecture proposed in [Biyong et al., 2020] is replacing this attention function with a *Sig-Attention* 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 [Biyong et al., 2020]) is used to ensure the output dimension is the same as the input dimension. The Sig-Attention mechanism is illustrated in Figure 17 below.

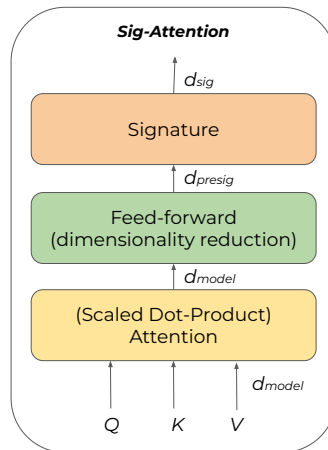


Figure 17: Sig-Attention mechanism—Part of Figure 2 in [Biyong et al., 2020]. Reproduced with permission. Colour Alterations. Licensed under a [Creative Commons Attribution 4.0 International License](#).

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 18 below which is part of Figure 2 in [Biyong et al., 2020]. The model takes the embedding of an input sequence as well as the queries, keys and values as input. As shown in Figure 18, 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 18: Additive Multi-Head Sig-Attention—Part of Figure 2 in [Biyong et al., 2020]. 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 11. This is a reproduction of Table 1 in [Biyong et al., 2020]

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 TABLETT 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 TABLETT 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 hjArtrytmen	1 table in the morning and a half table in the evening for heart rhythm	Cardiac- dysrhythmia	1.5	Standard
1-2 tablett 2 gÄnger dagligen	1-2 tablets 2 times daily	NA	3	Complex

Table 11: 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 [Biyong et al., 2020].

The following two approaches are proposed in [Biyong et al., 2020].

- Encode the Swedish text directly using Multilingual BERT (M-BERT) [Kenton and Toutanova, 2019], and
- Translate the prescriptions and then apply ClinicalBERT [Huang et al., 2019] that is pre-trained 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 19.

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 19 with the STE step removed) and against an LSTM model given by the same architecture as Figure 19, but with the STE step replaced by an LSTM step. The results are summarised in Table 12.

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 [Biyong et al., 2020] 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 [Biyong et al., 2020].

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 13.

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




Figure 19: Proposed multi-task learning architecture with Sig-Transformer Encoder (STE). Figure 4 in [Biyong et al., 2020]. Reproduced with permission. Colour Alterations. Licensed under a [Creative Commons Attribution 4.0 International License](#).

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 12: Performance comparisons for information extraction. Table 3 in [Biyong et al., 2020].

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 [Biyong et al., 2020] to create a relatively simple model whose performance is at least comparable with its competitors on the information extraction tasks considered in [Biyong et al., 2020].

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* (LHAR) 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 [Johansson, 1973]. 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

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 13: Performance for individual classes comparison; APPP means “As Per Previous Prescription” and NS means “Not Specified”. Table 6 in [Biyong et al., 2020].

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 [Vemulapalli et al., 2014] 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 [Yang et al., 2022], 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 [Jhuang et al., 2013, Fan et al., 2016], 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 21. 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

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

and

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

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 20.




Figure 20: 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 [Landmark Human Action Recognition Github](#).

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 [Yang et al., 2022]. 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 [Jhuang et al., 2013], SBU [Yun et al., 2012], Berkeley MHAD [Ofli et al., 2013] and NTURGB+D [Shahroudy et al., 2016]. 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 are cost effective depth cameras that are capable of providing reliable joint locations via real-time pose estimation algorithms [Shotton et al., 2013]. 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 2.5 for definitions of these transformations and Figure 21 for an illustration of the multi-delayed Lead-Lag transformations.




Figure 21: Illustration of multi-delayed lead-lag transformation. The dimension of lead-lag paths is d_{LL} . The delayed paths are padded with zeros to ensure a fixed length for each dimension. Figure 2 in [Yang et al., 2022]. 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 $n_{ST} := 3$ were found to be optimal in Section 6 of [Yang et al., 2022].

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 [Yang et al., 2022].

Illustrations of the extraction of the temporal features **T-J-PSF** and **T-S-PSF** are provided in Figures 23 and 24 respectively.

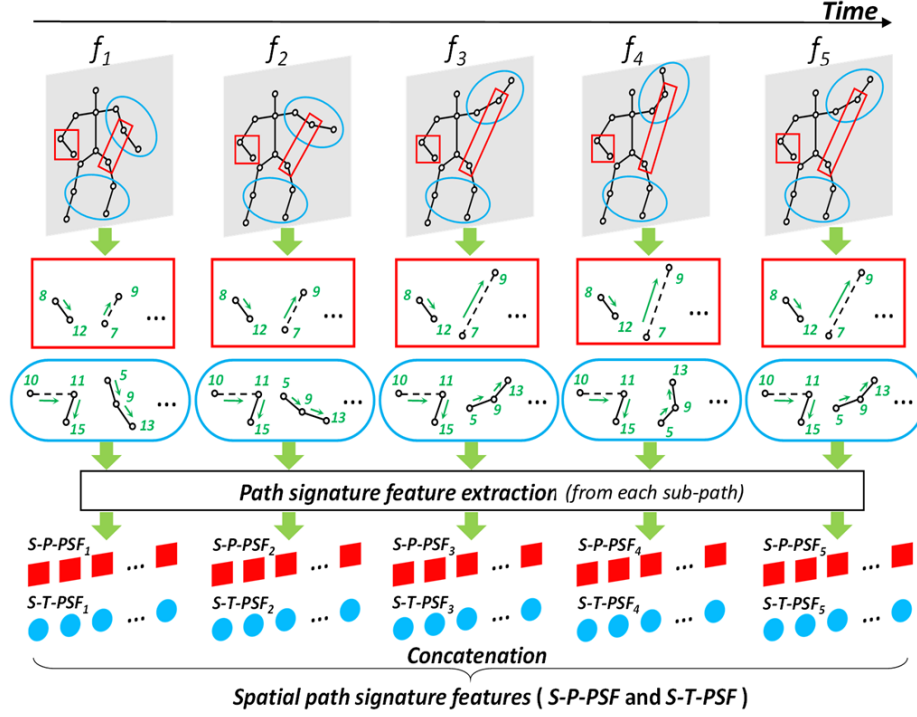


Figure 22: The illustration of spatial feature (**S-P-PSF** and **S-T-PSF**) extraction ($N=15$ in this figure). The red quadrangles denote the feature extraction of joint pairs, while the blue ellipses denote that of joint triples. All possible pairs and triples of joints are considered. Figure 3 in [Yang et al., 2022]. 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 [Yang et al., 2022] 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-PSF**) 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 [Wan et al., 2013], 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.




Figure 23: Illustration of temporal features extracted from the evolution of each corresponding joint (**T-J-PSF**). Figure 4 in [Yang et al., 2022]. Reproduced under license from Springer Nature (License Number 5425270874131).

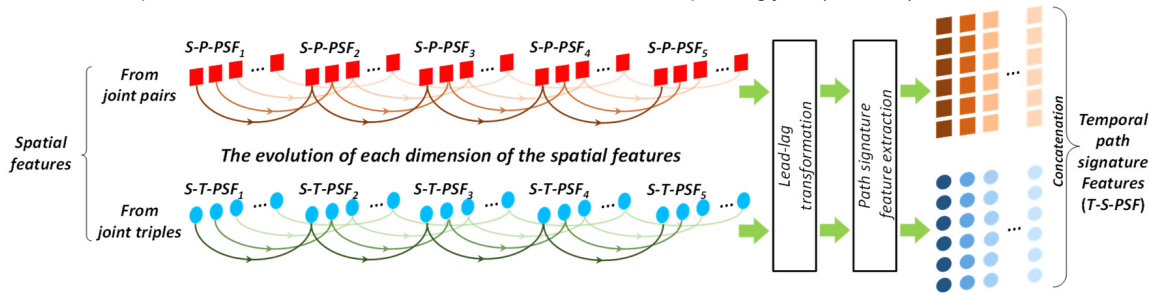


Figure 24: Illustration of temporal features extracted from the evolution of spatial context (**T-S-PSF**). Each dimension of the spatial features is treated equally and individually. Figure 5 in [Yang et al., 2022]. Reproduced under license from Springer Nature (License Number 5425270874131).

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 [Yang et al., 2022]. 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) [Fang et al., 2017] was used to obtain 17 estimated joints from the RGB videos. The signature method proposed in [Yang et al., 2022] is compared with the skeleton-based methods DT-FV [Wang and Schmid, 2013], HLPF [Jhuang et al., 2013], Novel HLPF [Fan et al., 2016], and P-CNN [Chéron et al., 2015], all trained using the evolution of these 17 joints, in [Yang et al., 2022]. As recorded in [Yang et al., 2022], the accuracy 80.4% achieved by the Path Signature method [Yang et al., 2022] is better than the accuracy achieved by any of the other methods (65.9% for DT-FV [Wang and Schmid, 2013], 74.6% for P-CNN [Chéron et al., 2015], 76.0% for HLPF [Jhuang et al., 2013], and 79.6% for Novel HLPF [Fan et al., 2016]).

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 [Yang et al., 2022] is compared with the skeleton-based methods Yun et al. [Yun et al., 2012], Ji et al. [Ji et al., 2014], CHARM [Li et al., 2015], HBRNN [Du et al., 2015], Deep HBRNN [Zhu et al., 2016], Co-occurrence [Zhu et al., 2016], STA-LSTM [Song et al., 2017], ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a], SkeletonNet [Ke et al., 2017], and GC-Attention-LSTM [Liu et al., 2017b], all trained using the evolution of the 30 joints, in [Yang et al., 2022]. As recored in Table 6 in [Yang et al., 2022], the accuracy of 96.8% for the Path Signature method [Yang et al., 2022] is better than the accuracy achieved by any of the other methods (80.3% for Yun et al. [Yun et al., 2012], 86.8% for Ji et al. [Ji et al., 2014], 83.9% for CHARM [Li et al., 2015], 80.4% for HBRNN [Du et al., 2015], 86.0% for Deep HBRNN [Zhu et al., 2016], 90.4% for Co-occurrence [Zhu et al., 2016], 91.5% for STA-LSTM [Song et al.,

2017], 93.3% for ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a], 93.5% for SkeletonNet [Ke et al., 2017], and 94.1% for GC-Attention-LSTM [Liu et al., 2017b]).

The joint locations are recorded precisely enough in the Berkely MHAD dataset to enable methods to achieve perfect accuracy. The signature method proposed in [Yang et al., 2022] is compared with the skeleton-based methods Vantigodi et al. [Vantigodi and Babu, 2013], Ofli et al. [Wang et al., 2013], Vantigodi et al. [Vantigodi and Radhakrishnan, 2014], Kapsouras et al. [Kapsouras and Nikolaidis, 2014], HBRNN [Du et al., 2015], and ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a] on the Berkeley MHAD dataset in [Yang et al., 2022]. As recorded in Table 7 in [Yang et al., 2022], the perfect accuracy of 100% achieved by the Path Signature method [Yang et al., 2022] matches the perfect accuracy of 100% achieved by both the HBRNN [Du et al., 2015] and SL-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a] methods, and is better than the accuracy achieved by the remaining methods (96.1% for Vantigodi et al. [Vantigodi and Babu, 2013], 95.4% for Ofli et al. [Wang et al., 2013], 97.6% for Vantigodi et al. [Vantigodi and Radhakrishnan, 2014], and 98.2% for Kapsouras et al. [Kapsouras and Nikolaidis, 2014]).

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 [Yang et al., 2022] is compared with the skeleton-based methods Dynamic Skeletons [Hu et al., 2015], HBRNN [Du et al., 2015], Part-aware LSTM [Shahroudy et al., 2016], ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a], STA-LSTM [Song et al., 2017], SkeletonNet [Ke et al., 2017], Joint Distance Maps [Li et al., 2017], and GC-Attention-LSTM [Liu et al., 2017b] in [Yang et al., 2022]. As recorded in Table 9 in [Yang et al., 2022], for the cross-subject task recognition, the accuracy of 78.3% achieved by the Path Signature method [Yang et al., 2022] is better than the accuracy achieved by any of the other methods (60.2% for Dynamic Skeletons [Hu et al., 2015], 59.1% for HBRNN [Du et al., 2015], 62.9% for Part-aware LSTM [Shahroudy et al., 2016], 69.2% for ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a], 73.4% for STA-LSTM [Song et al., 2017], 75.9% for SkeletonNet [Ke et al., 2017], 76.2% for Joint Distance Maps [Li et al., 2017], and 74.4% for GC-Attention-LSTM [Liu et al., 2017b]). Additionally, as recorded in Table 9 in [Yang et al., 2022], for the cross-view task recognition, the accuracy of 86.3% achieved by the Path Signature method [Yang et al., 2022] is better than the accuracy achieved by any of the other methods (65.2% for Dynamic Skeletons [Hu et al., 2015], 64.0% for HBRNN [Du et al., 2015], 70.3% for Part-aware LSTM [Shahroudy et al., 2016], 77.7% for ST-LSTM-Trust Gate [Liu et al., 2016, Liu et al., 2017a], 81.2% for STA-LSTM [Song et al., 2017], 81.2% for SkeletonNet [Ke et al., 2017], 82.3% for Joint Distance Maps [Li et al., 2017], and 82.8% for GC-Attention-LSTM [Liu et al., 2017b]).

The experiments conducted in [Yang et al., 2022] illustrate that, within the class of recognising actions via the evolution of connected segments between landmark locations, the Path Signature method proposed in [Yang et al., 2022], 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 [Yang et al., 2022]. It provides an introduction to the methodology of [Yang et al., 2022], 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 [Yang et al., 2022] (see Table 5 in [Yang et al., 2022]) 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. This notebook highlights the simplicity of the signature techniques; within 20 minutes it is possible to use the method proposed in [Yang et al., 2022] on a CPU to train a classifier that

achieves comparable accuracy with a specialised tailor-made models, see Table 5 in [Yang et al., 2022].

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 [Hill, 1986, Reichl, 2016, Schrödinger, 1989], to using high-resolution climatic data to predict overall end-of-year crop yields [Panda et al., 2010, Dahikar and Rode, 2014, You et al., 2017], and estimating mean-reversion parameters from observed financial market dynamics [Papavasiliou and Ladroue, 2011, Gatheral et al., 2018, Balvers et al., 2000].

In Section 6, an approach to distribution regression via the expected signature was outlined. This approach is implemented in [Lemercier et al., 2021], and we will provide a more detailed summary of their methodology within this section. Consider M input-output pairs $\left\{ \left\{ \mathbf{x}^{i,j} \right\}_{j=1}^{N_i}, y^i \right\}_{i=1}^M$ where each pair i is given by a target $y^i \in \mathbb{R}$ and a collection of N_i time-series $\mathbf{x}^{i,j} = \left\{ (t_1, x_1^{i,j}), \dots, (t_{l_{i,j}}, x_{l_{i,j}}^{i,j}) \right\}$ of lengths $l_{i,j} \in \mathbb{Z}_{\geq 1}$, time stamps $t_1 \leq \dots \leq t_{l_{i,j}}$ and values in a b -dimensional Banach space W . By using the method of transforming a stream of values into a stream of increments outlined in Subsection 2.3, we update each stream $\mathbf{x}^{i,j}$ to now be a stream of increments taking its values in a d -dimensional Banach space V . In particular, we must have that $d = 2(b+1) \in \mathbb{Z}_{\geq 2}$ and that, as a set, $V = W \times \mathbb{R}^{b+2}$.

Fix a compact interval $[a, b] \subset \mathbb{R}$. Following the notation used in [Lemercier et al., 2021], we let $\mathcal{C}([a, b], V)$ denote the set of paths resulting from the removal of all the tree-like (see [Hambly and Lyons, 2010]) paths from $\mathcal{V}^1([a, b], V)$. This has no practical impact on the following strategies [Lemercier et al., 2021]. Consider, for each $i \in \{1, \dots, M\}$ and $j \in \{1, \dots, N_i\}$ the contour $\Gamma_{\mathbf{x}^{i,j}}$ resulting from the concatenation of the entries of $\mathbf{x}^{i,j}$. Choose a parameterisation $x^{i,j} : [a, b] \rightarrow V$ of $\Gamma_{\mathbf{x}^{i,j}}$ such that $x^{i,j} \in \mathcal{C}([a, b], V)$.

After doing so, 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 represented as

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

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 two methods for distribution learning using the expected signature are proposed in [Lemercier et al., 2021]. 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

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

The *SES* method of [Lemercier et al., 2021] uses a truncation of $\mathcal{A}(Z)$, where $Z := \{\delta^i \mid i \in \{1, \dots, M\}\}$, 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 $K_0 \in \mathbb{Z}_{\geq 1}$ so that only the coordinate iterated integrals of order up to K_0 are used.

The second method introduced in [Lemercier et al., 2021] combines the expected signature with a Gaus-

sian kernel. Recall that the expected signature $\mathbb{S} : \mathcal{C}([a, b], V) \rightarrow T((V))$ is defined by

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

The KES model proposed in [Lemercier et al., 2021] 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 [Lemercier et al., 2021]).

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

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

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 [Salvi et al., 2021] (see Section 8) allows the inner products in (13.7) to be computed by a simple call to any numerical PDE solver of choice [Lemercier et al., 2021].

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 [Lemercier et al., 2021], the performance of the KES and SES models are compared with the performance of DeepSet method [Zaheer et al., 2017] 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 [Lemercier et al., 2021], with the full details appearing in [Smola et al., 2007], [Muandet et al., 2012] [Flaxman, 2015], [Szabó et al., 2016], [Law et al., 2018]. 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 [Cuturi et al., 2007], and the definition of all three may be found in Appendix B of [Lemercier et al., 2021]. 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 hyper-parameters 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 [Lemercier et al., 2021].

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 presented in Table 14.

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 14: 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 [Lemercier et al., 2021].

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 [Arribas et al., 2020]. Due to findings in [Gatheral et al., 2018], 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 [Lemercier et al., 2021]. 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 $\{\sigma_t^{i,j}\}_{j=1}^N$ of length 200. The results, presented in Table 15, illustrate that the models KES and SES out-perform the others, and that the performance of both KES and SES progressively with the number of samples N .

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

Table 15: 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 [Lemercier et al., 2021].

The final task considered in [Lemercier et al., 2021] is a crop yield prediction task. The goal is to predict 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. The results are presented in Table 16. Both the KES model and the SES model achieve larger improvements over the baseline than the other methods, with the SES model providing the best performance.

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

Table 16: Predicting wheat crop yield. Table 3 in [Lemercier et al., 2021].

14. Conformance and SigMahaKNN Method for Anomaly Detection

Anomaly detection is the task of determining whether a given observation is unusual compared to a corpus of observations that have been deemed to be usual. It arises in numerous fields such as medicine [Hauskrecht et al., 2013], financial fraud [Nian et al., 2016] and cybersecurity [Jones and Sielken, 2000]. 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 [Breunig et al., 2000], clustering [He et al., 2003], random forests [Liu et al., 2012], support vector machines [Amer et al., 2013], neural networks [Chalapathy and Chawla, 2019] and nearest neighbour-based approaches [Hautamaki et al., 2004, He and Wang, 2007, Sarmadi and Karamodin, 2020]. The article [Pimentel et al., 2014] surveys a wide-range of anomaly detection techniques.

Throughout the remainder of this section we consider the following generic problem framework. We assume that we have a corpus \mathcal{C} of observations that are all deemed to be normal. The aim is to use this corpus \mathcal{C} to determine whether a new observation *outside* of \mathcal{C} is an anomaly. In this setup we assume that the corpus \mathcal{C} is *not* polluted; no point within this given corpus \mathcal{C} should be considered anomalous. In particular, we do not consider the separate challenge of dealing with a polluted corpus where it is possible for points in the given corpus to be anomalous.

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 corpus \mathcal{C} . However, this viewpoint runs the risk of incorrectly deeming data points to be outliers.

For an example of incorrectly deeming a data point to be an outlier, suppose we are interested in the time taken by a person to run 100m. Assume we have an initial large corpus \mathcal{C} of people that contains at least one Olympic level sprinter. Let $\mu_{\mathcal{C}}$ denote the average time taken to run 100m taken over the corpus \mathcal{C} . If a new persons time is compared to $\mu_{\mathcal{C}}$, it is likely that an Olympic sprinter would be considered an outlier. But of course they are not an outlier since there is at least one *other* Olympic sprinter exhibiting similar performance amongst the corpus \mathcal{C} .

For an example of incorrectly deeming a data point to not be an outlier, assume we have a collection of points $\mathcal{C} \subset S^1$ where $S^1 := \{x \in \mathbb{R}^2 \mid |x| = 1\} \subset \mathbb{R}^2$ denotes the unit-circle in \mathbb{R}^2 with respect to the Euclidean distance on \mathbb{R}^2 . Assume that every point $(x, y) \in \mathbb{R}^2$ satisfies that the angle $\tan^{-1}(y/x)$ is rational. Provided the cardinality of \mathcal{C} is sufficiently large, the mean of \mathcal{C} will be the origin $(0, 0) \in \mathbb{R}^2$. Then suppose $(z, w) \in S^1$ with $\tan^{-1}(w/z) \in \mathbb{R} \setminus \mathbb{Q}$. The distance of (z, w) to the mean $(0, 0)$ is the same as the distance of *any* point $(x, y) \in \mathcal{C}$ to $(0, 0)$ (i.e. the distance is 1) and hence any distance-to-the-mean based anomaly detection procedure will not deem (z, w) an outlier. But (z, w) making an irrational angle to the positive x-axis means it differs from *every* point in the corpus \mathcal{C} , 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. This avoids the issue of incorrectly deeming a rare observation as anomalous. But, from this point-of-view, it is important to allow the corpus of observations that are deemed to be “normal” to be updated over time; after observing a rare event that is nevertheless *not* an anomaly, we want to add it to the “normal” corpus to ensure that future observations of this event will not be classified as an anomaly.

The *k-nearest neighbour* (k-NN) approach uses this idea by comparing an instance to its k nearest neigh-

bours 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 [Hautamaki et al., 2004]. 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 hyper-parameter 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* [Mahalanobis, 2018], providing a notion of distance between a point and a distribution, is a popular choice considered in the works [Si et al., 2020, Sarmadi and Karamodin, 2020], for example. The recent work [Shao et al., 2020] introduces the *variance norm* as a data-driven metric that offers a mathematically rigorous extension of the Mahalanobis distance. In particular, the variance norm coincides with the Mahalanobis distance when the sample covariance matrix is finite-dimensional and has full column rank [Shao et al., 2020]. However, unlike the Mahalanobis distance, the variance norm remains well-defined when these assumptions are *not* satisfied. As an example, multi-collinearity may often lead to data exhibiting rank-deficiency [Shao et al., 2020].

We turn our attention to the anomaly detection method proposed in [Shao et al., 2020] of using a notion of *conformance* to determine whether an entire stream is an anomalous object compared to a corpus of normal streams. A version of this work first appeared in 2020 [Perez Arribas et al., 2020]. Our presentation in this section focuses on the majorly updated version of the article [Shao et al., 2020] that appeared in December 2023.

Unlike earlier tasks considered in this article, anomaly detection is often *semi-supervised*. The data for training consists of a collection of streams of increments, with the collection denoted by $\Omega_{S(V)}$ to match our notation, that are normal objects. Only the normal objects are available during training, unlike the availability of two types of objects in binary classification problems. In fact, the anomaly detected might not even belong to the same class as the normal objects; their only defining characteristic is that they are different from all the objects in the corpus. A brief outline of the approach proposed in [Shao et al., 2020] is the following.

- i. Transform each stream $\mathbf{x} \in \Omega_{S(V)}$ to its path signature in $T((V))$.
- ii. Fit the data-driven variance norm to the resulting corpus of path signatures in $T((V))$.
- iii. Use the nearest neighbour algorithm in the tensor algebra $T((V))$ as a downstream metric-based anomaly detector.

Whilst the variance norm is considered on the tensor algebra $T((V))$ (or possibly on truncations of the tensor algebra), it actually makes sense in a general Banach space setting. Thus, to both match the framework assumed in [Shao et al., 2020] and avoid clashing with the notation we have adopted throughout this article, we fix a (possibly infinite dimensional) Banach space W and explain how the variance norm is defined with respect to a finite collection of elements $\mathcal{C} = \{w_1, \dots, w_n\} \subset W$ for some integer $n \in \mathbb{Z}_{\geq 1}$.

When $\dim(W) < \infty$ one can consider the Mahalanobis distance in this setting. To be more precise, we recall that if D is a distribution with mean μ and positive-definite covariance matrix S , then the Mahalanobis distance between a point y and the distribution D is defined to be

$$d_M(y, D) := \sqrt{(y - \mu)^T S^{-1} (y - \mu)} \quad (14.1)$$

where the superscript T denotes taking the transpose. If we take $D := \frac{1}{n} \sum_{i=1}^n \delta_{w_i}$ to be the empirical distribution summarising the corpus \mathcal{C} then $\mu = \mu_{\mathcal{C}} := \frac{1}{n} \sum_{i=1}^n w_i$ and $S = (X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})$ where X is the $(n \times d)$ matrix whose rows are the elements w_1, \dots, w_n with respect to some basis of W , and $\mathbf{1}_n := (1, \dots, 1) \in \mathbb{R}^n$. Then the Mahalanobis distance of $y \in W$ to the empirical distribution summarising the corpus \mathcal{C} is

$$d_M(y, X) := \sqrt{(y - \mu_{\mathcal{C}})^T [(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})]^{-1} (y - \mu_{\mathcal{C}})}. \quad (14.2)$$

Two limitations of (14.2) observed in [Shao et al., 2020] are:

- Computing a distance to the mean for anomaly detection may not be suitable for certain situations. The example of having the corpus \mathcal{C} uniformly distributed in the Euclidean unit circle $S^1 \subset \mathbb{R}^2$ discussed

at the start of this section illustrates this point. To provide a second example, let $S^2 \subset \mathbb{R}^3$ denote the Euclidean unit sphere in \mathbb{R}^3 , i.e. $S^2 := \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}$. Then suppose the corpus \mathcal{C} is uniformly distributed in $S^2_{z=0} := \{(x, y, z) \in S^2 \mid z = 0\}$. Provided the cardinality of the corpus \mathcal{C} is large enough, the mean will be the origin $(0, 0, 0) \in \mathbb{R}^3$. Hence in terms of the distance to the mean, the point $(0, 0, 1) \in S^2$ is not anomalous compared to the corpus \mathcal{C} . However, it evidently should be classified as an anomaly since it differs from every point in the corpus \mathcal{C} .

- It is unclear how to proceed when the empirical distribution covariance matrix $(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})$ is singular. This is often the case when, for example, the streamed data exhibits multi-colinearity between the channels.

These limitations are addressed by the variance norm introduced in [Shao et al., 2020]. We now define this notion in the setting of the corpus $\mathcal{C} = \{w_1, \dots, w_n\} \subset W$, and stress that there is no assumption made on the dimension of W . Let $\mu_{\mathcal{C}} := \frac{1}{n} \sum_{i=1}^n w_i$ be the mean of the corpus \mathcal{C} . We first define a bilinear form $\mathbf{q} : W^* \times W^* \rightarrow \mathbb{R}$ by setting, for $\sigma, \varphi \in W^*$,

$$\mathbf{q}(\sigma, \varphi) := \sum_{i=1}^n \sigma(w_i - \mu_{\mathcal{C}}) \varphi(w_i - \mu_{\mathcal{C}}). \quad (14.3)$$

Subsequently define a map $\mathbf{p} : W \times W \rightarrow \mathbb{R}$ by setting, for $w, u \in W$,

$$\mathbf{p}(w, u) := \sup \left\{ \sigma(w - \mu_{\mathcal{C}}) \sigma(u - \mu_{\mathcal{C}}) \mid \sigma \in W^* \text{ with } \mathbf{q}(\sigma, \sigma) \leq 1 \right\}. \quad (14.4)$$

Then the *variance norm* of $w \in W$ is defined to be

$$\|w\|_{\text{var}} := \sqrt{\mathbf{p}(w, w)}. \quad (14.5)$$

When W is finite dimensional, the following key properties of the variance norm $\|\cdot\|_{\text{var}}$ are established in Theorem 3.1 in [Shao et al., 2020].

Theorem 14.1 (Properties of $\|\cdot\|_{\text{var}}$ when $\dim(W) < \infty$; Theorem 3.1 in [Shao et al., 2020]). *Assume W is a finite dimensional real Banach space, $n \in \mathbb{Z}_{\geq 1}$ and $\mathcal{C} = \{w_1, \dots, w_n\} \subset W$. Fix a basis for W and let X be the matrix whose rows are the elements w_1, \dots, w_n expressed with respect to this basis. Let $\mathbf{1}_n := (1, \dots, 1) \in \mathbb{R}^n$ and $\mu_{\mathcal{C}} := \frac{1}{n} \sum_{i=1}^n w_i$ be the mean of the corpus \mathcal{C} . Then the following properties are true.*

- (A) *If $w \in \text{Span}(\mathcal{C})$ then $\|w\|_{\text{var}}^2 = \mathbf{p}(w, w) = (w - \mu_{\mathcal{C}})^T [(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})]^{\dagger} (w - \mu_{\mathcal{C}})$ where the superscript \dagger is used to denote the Moore-Penrose pseudo-inverse.*
- (B) *If $w \notin \text{Span}(\mathcal{C})$ then $\|w\|_{\text{var}}^2 = \mathbf{p}(w, w) = \infty$.*

We stress that the choice of a basis of W in Theorem 14.1 is purely for the purposes of expressing the matrix X . The result itself is *not* dependent on the choice of basis; one could select any other basis and obtain the same result, only now with the matrix X expressed with respect to the new basis of W instead.

Using the variance norm we define the *conformance distance* of an element $w \in W$ to the corpus \mathcal{C} is defined to be (see Definition 3.2 in [Shao et al., 2020])

$$d_{\mathcal{C}}(w, \mathcal{C}) := \min_{i \in \{1, \dots, n\}} \|w - w_i\|_{\text{var}}^2. \quad (14.6)$$

Thus in the case that W is finite dimensional we have, via Theorem 14.1, for every $w \in W$ that either $d_{\mathcal{C}}(w, \mathcal{C}) = \infty$ if $w \notin \text{Span}(\mathcal{C})$, or

$$d_{\mathcal{C}}(w, \mathcal{C}) = (w - w_i - \mu_{\mathcal{C}})^T [(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})]^{\dagger} (w - w_i - \mu_{\mathcal{C}}) \quad (14.7)$$

for some $i \in \{1, \dots, n\}$ if $w \in \text{Span}(\mathcal{C})$. It follows from (14.7) that if the matrix X has full column rank and $w \in \text{Span}(\mathcal{C})$ then the conformance distance is equal to the nearest neighbour Mahalanobis distance [Shao et al., 2020].

Allowing the conformance distance to be infinity in some cases is a distinction from most Mahalanobis distance based anomaly detection methods [Shao et al., 2020]. But this provides the benefit of avoiding an insensitivity introduced by the use of the pseudo-inverse of the matrix $(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})$. As explained in section 3.2.2 of [Shao et al., 2020], the quantity $(w - w_i - \mu_{\mathcal{C}})^T [(X - \mathbf{1}_n \mu_{\mathcal{C}})^T (X - \mathbf{1}_n \mu_{\mathcal{C}})]^\dagger (w - w_i - \mu_{\mathcal{C}})$ remains invariant under translation of w by a particular class of elements $z \in W$. Loosely, the class of elements is determined via the SVD of the matrix X ; see section 3.2.2 in [Shao et al., 2020] for full details.

This is problematic from an anomaly detection perspective since the class of elements $z \in W$ for which this quantity is invariant under $w \mapsto w + z$ enables one to map $w \in \text{Span}(\mathcal{C})$ to an element $w + z \notin \text{Span}(\mathcal{C})$. Being outside $\text{Span}(\mathcal{C})$ ought to make $w + z$ an anomaly; but the invariance means it will not be marked as an outlier by a method based solely on the pseudo-inverse. However, the conformance distance does mark it as an exceptional point by returning the value ∞ ; despite such elements determining a different type of anomaly, the conformance distance remains able to detect them [Shao et al., 2020].

We now illustrate how the variance norm and conformance distance are used in [Shao et al., 2020] to define the SigMahaKNN method for detecting anomalies in streamed data. In the notation used in this article, the SigMahaKNN method is developed in the setting that $V := \mathbb{R}^d$ for some integer $d \in \mathbb{Z}_{\geq 1}$. That is, for an integer $M \in \mathbb{Z}_{\geq 1}$, we have a collection of streams of increments $\Omega_{\mathcal{S}(\mathbb{R}^d)} = \{\mathbf{x}_1, \dots, \mathbf{x}_M\} \subset \mathcal{S}(\mathbb{R}^d)$. The SigMahaKNN method is designed to classify whether a stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^d) \setminus \Omega_{\mathcal{S}(\mathbb{R}^d)}$ is an outlier compared to the collection $\Omega_{\mathcal{S}(\mathbb{R}^d)}$.

We first pick an integer $N \in \mathbb{Z}_{\geq 1}$ and consider the corpus $\mathcal{C} := \{S^{(N)}(\mathbf{x}) \mid \mathbf{x} \in \Omega_{\mathcal{S}(\mathbb{R}^d)}\} \subset T^{(N)}(\mathbb{R}^d)$ of elements in the truncated tensor algebra $T^{(N)}(\mathbb{R}^d)$ obtained by taking the truncation to depth N of the signature of each stream $\mathbf{x} \in \Omega_{\mathcal{S}(\mathbb{R}^d)}$ (cf. Section 3). Recalling Subsection 2.25, we can view the corpus \mathcal{C} as a collection of M points in the Euclidean space $\mathbb{R}^{m(d,N)}$ for

$$m = m(d, N) := \sum_{j=0}^N d^j = \begin{cases} N+1 & \text{if } d = 1 \\ \frac{d^{N+1}-1}{d-1} & \text{if } d \geq 2. \end{cases} \quad (14.8)$$

For notational convenience, we let $\mathcal{C} = \{z_1, \dots, z_M\} \subset \mathbb{R}^M$. Define an empirical measure $\delta_{\mathcal{C}} := \frac{1}{M} \sum_{j=1}^M \delta_{z_j}$ summarising \mathcal{C} (cf. Section 13), and subsequently let $\mu_{\mathcal{C}} := \frac{1}{M} \sum_{j=1}^M z_j$ denote its mean. Finally, to match the notation adopted in [Shao et al., 2020], we let $\|\cdot\|_{\text{Sig}^N(\mathcal{C})}$ denote the conformance distance $d_{\mathcal{C}}(\cdot, \mathcal{C})$ defined by (14.6) for the corpus $\mathcal{C} \subset \mathbb{R}^m$.

It is $\|\cdot\|_{\text{Sig}^N(\mathcal{C})}$ that determines the anomaly score for the SigMahaKNN algorithm; streams with conformance distance higher than a user-determined threshold are anomalies. That is, suppose we determine an appropriate threshold $a > 0$. Then let $\mathbf{x} \in \mathcal{S}(\mathbb{R}^d)$ be a stream that is not in the original collection $\Omega_{\mathcal{S}(\mathbb{R}^d)}$. The SigMahaKNN method determines whether \mathbf{x} is an anomaly compared to the collection $\Omega_{\mathcal{S}(\mathbb{R}^d)}$ as follows.

- (A) Compute $S^{(N)}(\mathbf{x}) \in T^{(N)}(\mathbb{R}^d) \cong \mathbb{R}^m$, the truncated to depth N signature of the stream \mathbf{x} .
- (B) Compute $\|S^{(N)}(\mathbf{x})\|_{\text{Sig}^N(\mathcal{C})}$, the conformance distance of $S^{(N)}(\mathbf{x})$ relative to the corpus \mathcal{C} .
- (C) If $\|S^{(N)}(\mathbf{x})\|_{\text{Sig}^N(\mathcal{C})} > a$ then the stream \mathbf{x} is classified as an anomaly; otherwise the stream \mathbf{x} is classified as normal.

An obvious question is how the threshold value $a > 0$ should be determined. The authors propose using the nearest neighbour distance to set an appropriate threshold in [Shao et al., 2020]. This is done as follows.

- (A) Randomly split $\mathcal{C} = \mathcal{C}_1 \sqcup \mathcal{C}_2$ into two equal-sized parts \mathcal{C}_1 and \mathcal{C}_2 .
- (B) Compute the empirical cdf of the nearest neighbour distance for \mathcal{C}_1 using \mathcal{C}_2 as the corpus.

(C) Choose an appropriate tail quantile in the empirical cdf to set the threshold value $a > 0$.

The choice of the depth $N \in \mathbb{Z}_{\geq 1}$ to which the signature is truncated determines a parameter that may be varied. 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.

The following theoretical properties of the conformance distance $\|\cdot\|_{\text{Sig}^N(\mathcal{C})}$ as the truncation level $N \in \mathbb{Z}_{\geq 1}$ varies are established in [Shao et al., 2020]. The first establishes that for a fixed stream $\mathbf{x} \in \mathcal{S}(\mathbb{R}^d)$, the mapping $N \mapsto \|S^{(N)}(\mathbf{x})\|_{\text{Sig}^N(\mathcal{C})}$ is monotonically increasing (see Proposition 4.1 in [Shao et al., 2020]). Hence the extent to which a stream \mathbf{x} is considered an outlier compared to the original collection $\Omega_{\mathcal{S}(\mathbb{R}^d)}$ increases as the resolution N at which we examine the streams increases.

The second establishes that if $\mathbf{x} \notin \Omega_{\mathcal{S}(\mathbb{R}^d)}$ then there exists an integer $N_* \in \mathbb{Z}_{\geq 1}$ such that for every integer $N \in \mathbb{Z}$ with $N \geq N_*$ the conformance distance of $S^{(N)}(\mathbf{x})$ with respect to the corpus \mathcal{C} is ∞ (see Proposition 4.2 in [Shao et al., 2020]). Thus if a stream \mathbf{x} is *not* in the original collection $\Omega_{\mathcal{S}(\mathbb{R}^d)}$ of streams, we are guaranteed that it will be classified as an outlier provided we examine the streams at a sufficiently high resolution.

Four major benefits of the SigMahaKNN method proposed in [Shao et al., 2020] are:

- It is *dimensionless* in the sense that it is independent of the unit of measurement of the streams. Moreover it is invariant to time-reparameterisation and concatenation of streams. These are direct consequence of the invariance of the signature covered in Theorem 3.1 in this article. Particular theoretical results establishing these property for the SigMahaKNN method may be found in Theorems 4.1, 4.2, and 4.3 in [Shao et al., 2020].
- It is a data-driven extension of Mahalanobis distance based models that continues to make sense even when the sample covariance matrix is singular. In particular, it can deal with data exhibiting multi-collinearity between channels.
- No distribution assumptions are required for the original collection of streams $\Omega_{\mathcal{S}(\mathbb{R}^d)} \subset \mathcal{S}(\mathbb{R}^d)$. The method can be applied to any finite collection of streams in $\mathcal{S}(\mathbb{R}^d)$.
- Different types of anomalies that are missed by Mahalanobis distance based models are detected by the SigMahaKNN model.

For the remainder of this section we turn our attention to the numerical performance of the SigMahaKNN method as reported in section 5 of [Shao et al., 2020]. A detailed algorithm for the implementation of the SigMahaKNN method is provided in algorithm 1 in [Shao et al., 2020]. In particular, it includes consideration of numerical rank deficiency issues involved in the computation of the Moore-Penrose pseudo-inverse via SVD factorisation that we have omitted.

For the two example problems we present here, the SigMahaKNN method is compared with *isolation forest* [Liu et al., 2008] and the *local outlier factor method* [Breunig et al., 2000]. In order to use these well-known anomaly detection methods on streamed data, the authors of [Shao et al., 2020] consider them with either the moment features (mean and covariance of different channels of the stream), of the signature features. This leads to four baselines: **IF-M** (Isolation Forest with Moment features), **IF-S** (Isolation Forest with Signature features), **LOF-M** (Local Outlier Factor with Moment features), and **LOF-S** (Local Outlier Factor with Signature features).

The SigMahaKNN method is evaluated using the PenDigitis-orig data set [Alpaydin and Alimoglu, 1996], consisting of 10992 instances of handwritten digits captured from 44 subjects, in section 5 of [Shao et al., 2020]. Each instance is represented as a two-dimensional stream, based on sampling of the pen position.

Given a digit k , a set of normal data $\mathcal{I}_{\text{Normal}}$ is defined to be the set of instances representing the digit k . The subset of $\mathcal{I}_{\text{Normal}}$ labelled as ‘training’ by the annotators is taken for the collection $\Omega_{\mathcal{S}(\mathbb{R}^2)}$ of normal streams. Further \mathcal{V} is taken to be the set of instances labelled as ‘testing’ by the annotators, which results in $\#(\mathcal{V}) = 3498$. Finally, $\mathcal{I}_{\text{Anomaly}}$ is the subset of \mathcal{V} not representing digit k . On average, taken over the consideration of all possible digits k , one has the cardinalities $\#(\Omega_{\mathcal{S}(\mathbb{R}^2)}) = 749.4$ and $\#(\mathcal{I}_{\text{Anomaly}}) =$

3148.2. Min-Max normalisation is applied to each individual stream since each digit class is invariant to translation and scaling.

Truncated signatures of orders $N \in \{1, \dots, 5\}$ are considered without any augmentation. The results, based on aggregating conformance values across the set of possible digits, are summarised in Table 17. Once the truncation level is at least 2 the SigMahaKNN out-performs all the baseline comparison models.

	$N = 1$	$N = 2$	$N = 3$	$N = 4$	$N = 5$
SigMahaKNN	0.870	0.942	0.948	0.954	0.956
IF-M	-	-	-	-	0.618
IF-S	0.888	0.931	0.916	0.875	0.834
LOF-M	-	-	-	-	0.514
LOF-s	0.563	0.584	0.582	0.582	0.582

Table 17: Handwritten digits data: performance quantified using ROC AUC in response to signature order N . Bootstrapped standard errors based on 10^4 samples are around 0.003. Table 1 in [Shao et al., 2020]

The second numerical example from [Shao et al., 2020] that we cover 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. This dataset, collected by the US Coast Guard in January 2017, contains a total of 31 884 021 geographical positions recorded for 6282 distinct vessel identifiers. The authors of [Shao et al., 2020] consider the stream of timestamped latitude/longitude position data associated with each vessel.

Vessel data with invalid identifiers of invalid length information were discarded. Streams are compressed to only contain points that are at least 10m away from the previous position, and, to help ensure that streams are faithful representations of ship movement, 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.

A finite collection of streams is obtained from 607 vessels whose sub-streams total between 10111 ($D = 32\text{km}$) and 104369 ($D = 4\text{km}$). A subset of normal instances used for testing is obtained from 607 vessels whose sub-streams total between 11254 ($D = 32\text{km}$) and 114071 ($D = 4\text{km}$). Finally, a subset of anomalous instances is obtained from 997 vessels whose sub-streams total between 8890 ($D = 32\text{km}$) and 123237 ($D = 4\text{km}$).

Imbalance in the number of sub-streams is accounted for by using, for each of the aforementioned three subsets, a weighted sample of 5000 instances. After computing the sub-streams and transforming them as outlined above, Min-Max normalisation is applied. The difference between successive timestamps is augmented into the streams as a new channel to account for velocity (cf. Time Augmentation in Subsection 2.5). Signatures truncated to depth $N = 3$ are considered. For baseline comparison models, each sub-stream is also summarised by estimating its component-wise mean and covariance, retaining the upper triangular part of the covariance matrix. The resulting set of features has dimensionality $\frac{1}{2}(n^2 + 3n)$, and these features are provided as input to an isolation forest [Liu et al., 2008]. The isolation forest is trained using 100 trees and for each tree in the ensemble using 256 samples represented by a single random feature.

Table 18 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 . Signature conformance yields higher ROC-AUC scores than the baseline models for the majority of the parameter combinations. The maximum ROC-AUC score (highlighted in bold in Table 18) 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 10 percentage points gain compared to the best-performing baselines parameter combination for $D = 32\text{km}$. This best-performing baseline parameter combination for $D = 32\text{km}$ is achieved by the IF-M model using a combination of time differences and invisibility reset transformations, achieving a ROC-AUC score of 0.828.

Model	Lead-Lag	Time-Diff	Time-inv	D = 4km	D = 8km	D = 16km	D = 32km
SigMahaKNN	No	No	No	0.723	0.706	0.705	0.740
IF-M	No	No	No	0.714	0.712	0.727	0.727
IF-S	No	No	No	0.627	0.623	0.645	0.660
LOF-M	No	No	No	0.543	0.542	0.522	0.513
LOF-S	No	No	No	0.484	0.500	0.491	0.492
SigMahaKNN	No	No	Yes	0.776	0.789	0.785	0.805
IF-M	No	No	Yes	0.781	0.785	0.776	0.790
IF-S	No	No	Yes	0.686	0.701	0.715	0.718
LOF-M	No	No	Yes	0.555	0.585	0.564	0.535
LOF-S	No	No	Yes	0.565	0.572	0.562	0.520
SigMahaKNN	No	Yes	No	0.810	0.813	0.818	0.848
IF-M	No	Yes	No	0.767	0.772	0.786	0.804
IF-S	No	Yes	No	0.731	0.714	0.737	0.771
LOF-M	No	Yes	No	0.547	0.543	0.526	0.520
LOF-S	No	Yes	No	0.511	0.505	0.493	0.494
SigMahaKNN	No	Yes	Yes	0.839	0.860	0.863	0.879
IF-M	No	Yes	Yes	0.830	0.823	0.831	0.828
IF-S	No	Yes	Yes	0.777	0.784	0.789	0.808
LOF-M	No	Yes	Yes	0.559	0.589	0.572	0.545
LOF-S	No	Yes	Yes	0.565	0.572	0.561	0.520
SigMahaKNN	Yes	No	No	0.811	0.835	0.824	0.837
IF-M	Yes	No	No	0.696	0.704	0.711	0.724
IF-S	Yes	No	No	0.617	0.596	0.634	0.668
LOF-M	Yes	No	No	0.543	0.541	0.522	0.513
LOF-S	Yes	No	No	0.484	0.500	0.491	0.493
SigMahaKNN	Yes	No	Yes	0.812	0.835	0.833	0.855
IF-M	Yes	No	Yes	0.758	0.759	0.767	0.773
IF-S	Yes	No	Yes	0.692	0.701	0.691	0.713
LOF-M	Yes	No	Yes	0.566	0.556	0.542	0.533
LOF-S	Yes	No	Yes	0.564	0.569	0.558	0.518
SigMahaKNN	Yes	Yes	No	0.845	0.861	0.862	0.877
IF-M	Yes	Yes	No	0.747	0.763	0.780	0.785
IF-S	Yes	Yes	No	0.725	0.692	0.716	0.757
LOF-M	Yes	Yes	No	0.547	0.543	0.526	0.520
LOF-S	Yes	Yes	No	0.530	0.533	0.553	0.600
SigMahaKNN	Yes	Yes	Yes	0.848	0.863	0.870	0.891
IF-M	Yes	Yes	Yes	0.811	0.813	0.809	0.823
IF-S	Yes	Yes	Yes	0.779	0.782	0.801	0.823
LOF-M	Yes	Yes	Yes	0.572	0.563	0.554	0.544
LOF-S	Yes	Yes	Yes	0.574	0.588	0.589	0.584

Table 18: SigMahaKNN, IF-M, IF-S, LOF-M, and LOF-S on Marine vessel traffic data: performance quantified using ROC AUC. Best across all transformations and models are in bold. Combination of Tables 2, 3, 4, 5, and 6 in [Shao et al., 2020]

Subsequent to the appearance of [Shao et al., 2020], the outlier detection framework of the SigMahaKNN method is considered for anomaly detection on radio astronomy data in [Arrubarrena et al., 2024]. In particular, the authors of [Arrubarrena et al., 2024] develop the SigNova method, inspired by the SigMahaKNN framework, that can detect much fainter *radio frequency interference* (RFI) than the two main frameworks SSINS [Wilensky et al., 2019] and AOFlagger [Offringa et al., 2012, Offringa et al., 2015] are capable of detecting. The SigNova method implements all the steps of the SigMahaKNN framework with an additional application of the `pysegments` segmentation algorithm [Arrubarrena et al., 2024].

In [Cass et al., 2024] the authors propose a novel framework for Mahalanobis-type anomaly detection on any Banach space (including, in particular, any Hilbert space) that is based upon the generalised notion of variance norm introduced in [Shao et al., 2020] and ideas from Cameron-Martin spaces (see e.g. [Lifshits and Lifshits, 2012, Bogachev, 1998]). In the Hilbert space setting, it is established that the variance norm depends

solely on the given inner product, and hence that the kernelised Mahalanobis distance can be recovered by working on *Reproducing Kernel Hilbert Spaces* (RKHS) [Cass et al., 2024]. This underpins the authors' introduction of the *Kernelised Nearest-Neighbour Mahalanobis* distance for semi-supervised anomaly detection in [Cass et al., 2024]. Both theoretical justifications and empirical validations of this approach can be found in [Cass et al., 2024].

The paper [Akyildirim et al., 2022] considers the use of path signature techniques for anomaly detection with the aim of detecting market manipulation attempts from financial data. There are two main reasons we delay discussing this approach until Section 15. The first is that the authors of [Akyildirim et al., 2022] make use of “off-the-shelf” anomaly detection algorithms without alteration; this is quite distinct in spirit from the approach in [Shao et al., 2020] of introducing a novel measure of similarity related to a reformulation of what being an anomaly *should* mean. The second is that a major novelty of the approach proposed in [Akyildirim et al., 2022] is the use of *randomised signatures*, which are noteworthy enough to warrant their own dedicated section.

15. Randomised Signature

In this section we discuss *randomised signatures* as presented in, for example, [Akyildirim et al., 2022]. Loosely speaking, randomised signatures exploit the notion of *reservoir computing* for modelling input-output systems to provide more concrete realisations (or, algebraically speaking, representations) of the abstract Tensor algebra (cf. Subsection 2.7). The close link between path signatures and CDEs discussed in Section 4 is central to this approach. We begin our presentation with a brief introduction to reservoir computing.

Fix non-negative integers $n, k, e \in \mathbb{Z}_{\geq 1}$. *Reservoir Computers* (RCs) are state-space transformations given by recurrent neural networks (RNNs) determined by two maps; namely, a *reservoir* (or *state*) map $F : \mathbb{R}^n \times \mathbb{R}^e \rightarrow \mathbb{R}^n$, and a *readout* map $h : \mathbb{R}^n \rightarrow \mathbb{R}^k$. The integer $n \in \mathbb{Z}_{\geq 1}$ is the number of virtual neurons of the system. RCs transform (or filter) a discrete-time input $z = (z_m)_{m \in \mathbb{Z}} \subset \mathbb{R}^e$ into an output signal $y = (y_m)_{m \in \mathbb{Z}} \subset \mathbb{R}^k$ using the state-space transformation given, for $t \in \mathbb{Z}$, by

$$y_t := h(x_t) \in \mathbb{R}^k \quad \text{for} \quad x_t := F(x_{t-1}, z_t) \in \mathbb{R}^n. \quad (15.1)$$

The RC satisfies the *echo state property* (ESP) if any input signal $(z_m)_{m \in \mathbb{Z}} \subset \mathbb{R}^e$ results in a unique output under the transformation detailed in (15.1). Typically, the reservoir map F does *not* depend on the specific dynamics, and instead provides a collection (or reservoir) of features. The readout map h is the object that is actually trained to use the features provided by the reservoir map F to match the observed collection of outputs sufficiently well.

The RC approach enables one to translate infinite dimensional problems regarding filters to analogous questions related to the reservoir and readout maps that generate the filter and are defined on finite dimensional spaces. An in-depth discussion of reservoir computing covering numerous important theoretical results can be found in [Grigoryeva and Ortega, 2018], [Cuchiero et al., 2021a], and the references there in. A particularly noteworthy result is the universality property established for a particular class of RCs in [Grigoryeva and Ortega, 2018]. In [Grigoryeva and Ortega, 2018] the authors establish that the class of RCs for which the reservoir map $F : \mathbb{R}^n \times \mathbb{R}^e \rightarrow \mathbb{R}^n$ is given by $F(x, y) := \sigma(\mathbf{A}x + \mathbf{B}y + \xi)$, for real matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times e}$, a vector $\xi \in \mathbb{R}^n$, and a sigmoid function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ applied component wise, and the readout map $h : \mathbb{R}^n \rightarrow \mathbb{R}^k$ is given by $h(x) := \mathbf{W}x$, for a real matrix $\mathbf{W} \in \mathbb{R}^{k \times n}$, are universal uniform approximants for a large class of discrete-time filters; see Theorem 4.1 in [Grigoryeva and Ortega, 2018] for the full detailed statement.

Reservoir computing offers an approach to the task of describing the solution trajectories of CDEs without having access to the vector fields. To be more precise, fix non-negative integers $d, l \in \mathbb{Z}_{\geq 1}$, fix $T > 0$, and (for simplicity) suppose that $X : [0, T] \rightarrow \mathbb{R}^d$ is a continuous path of bounded variation, i.e. $X \in \mathcal{V}^1([0, T]; \mathbb{R}^d)$. Further suppose that $f : \mathbb{R}^l \rightarrow \mathbf{L}(\mathbb{R}^d, \mathbb{R}^l)$ is continuous. Then, for a given $x_0 \in \mathbb{R}^l$, we may consider trying to find a path $z : [0, T] \rightarrow \mathbb{R}^l$ solving the following CDE driven by X :

$$z_0 = x_0 \quad \text{and} \quad dz_t = f(z_t)dX_t \quad (15.2)$$

for every $t \in (0, T]$. In Section 4 we observed that if the CDE (15.2) is linear in the sense that $f(z_t)dX_t = B(dX_t)z_t$ for a bounded linear map $B : \mathbb{R}^d \rightarrow \mathbf{L}(\mathbb{R}^l, \mathbb{R}^l)$, then the solution to (15.2) is given by a linear combination of the terms of the signature of the path X . To be more precise, in this case the solution to (15.2) is given by (cf. (4.3))

$$z_t = \left(\sum_{i=0}^{\infty} B^{\otimes i} \left(S_{0,t}^i(X) \right) \right) x_0 \quad (15.3)$$

Evidently, for linear CDEs, one approach is to take the (truncated) signature as the reservoir map, and subsequently take the readout map to be a trainable linear map. This strategy is completely analogous to the generic strategy followed throughout this survey article.

However, as noted in [Cuchiero et al., 2021b], the difficulty in calculating the precise values of the reservoir given by the signature of a control path $X \in \mathcal{V}^1([0, T]; \mathbb{R}^d)$ is in sharp contrast to the usual setting for reservoir computing in which reservoirs are easy to evaluate. In [Cuchiero et al., 2021b] the authors overcome this issue by introducing the *randomised signature*. The information of the signature is compressed via a random projection to a lower dimensional Euclidean space. A key aspect is that the image of this projection can be well-approximated by a dynamical system on the lower dimensional Euclidean space with random characteristics. In order to provide a more detailed explanation of the randomised signature we additionally make use of the presentation provided in [Akyildirim et al., 2022].

Coarsely, the idea of randomised signatures is to provide a more concrete realisation of the tensor algebra $T((\mathbb{R}^{d+1}))$ via mapping the canonical basis e_0, \dots, e_d of \mathbb{R}^{d+1} to a collection of informative vector fields $\mathbb{R}^k \rightarrow \mathbb{R}^k$ for some moderately sized integer $k \in \mathbb{Z}_{\geq 1}$. The switch from d to $d+1$ reflects that path X is replaced by its time-augmented variant where the running time is recorded as a new additional zeroth component. The choice of the form of vector fields is partly motivated by both the aim to model a random projection of the signature, and the universality results established in [Grigoryeva and Ortega, 2018] and [Cuchiero et al., 2021a].

To specify the class of vector fields $\mathbb{R}^k \rightarrow \mathbb{R}^k$ considered, fix a choice of a bounded real analytic activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, and, for each $j \in \{0, \dots, d\}$, a choice of $k \times k$ matrix $\mathbf{A}_j \in \mathbb{R}^{k \times k}$ and vector $b_j \in \mathbb{R}^k$. Then one may consider the map from \mathbb{R}^{d+1} to the set of vector fields on \mathbb{R}^k given by mapping, for each $j \in \{0, \dots, d\}$, the basis element e_j to the vector field $f_j : \mathbb{R}^k \rightarrow \mathbb{R}^k$ defined, for $x \in \mathbb{R}^k$, by

$$f_j(x) := \sigma(\mathbf{A}_j x + b_j). \quad (15.4)$$

In (15.4), the activation function σ is applied component-wise. Following the notation used in Definition 2.12 in [Akyildirim et al., 2022], the solution $\text{RS} : [0, T] \rightarrow \mathbb{R}^k$ to the linear CDE

$$\text{RS}_0 \in \mathbb{R}^k \quad \text{with} \quad d \text{RS}_t = \sum_{j=0}^d f_j(\text{RS}_t) dX_t^j = \sum_{j=0}^d \sigma(\mathbf{A}_j \text{RS}_t + b_j) dX_t^j \quad (15.5)$$

for every $t \in (0, T]$ is called *randomised signature*.

Theoretical results in both [Cuchiero et al., 2021b] and [Akyildirim et al., 2022] establish that, under certain assumptions, randomised signatures are universal approximants in the space of continuous functions. The approach in [Cuchiero et al., 2021b] is inspired by reservoir computing techniques, culminating in Theorem 3.3 which makes use of the *Johnson–Lindenstrauss* lemma to establish that randomised signatures are as expressive as the signature itself in a quantified sense. In a similar spirit, Theorem 2.13 in [Akyildirim et al., 2022] gives a more direct argument to verify that randomised signatures inherit the universality of the full signature provided the representation from the free algebra generated by e_0, \dots, e_d to the algebra of differential operators on \mathbb{R}^k defined in (15.4) is injective. It is additionally required in Theorem 2.13 in [Akyildirim et al., 2022] that $b_0, \dots, b_d, \mathbf{A}_0, \dots, \mathbf{A}_d$ are independent samples of a probability law absolutely continuous with respect to the Lebesgue measure.

We end this section with an illustration of the use of randomised signature for anomaly detection presented in [Akyildirim et al., 2022]. In particular, the task of detecting market fraud by identifying so-called *Pump-and-Dump* (PD) attempts within various cryptocurrencies is considered in [Akyildirim et al., 2022]. The lack of regulation of cryptocurrency brokers has resulted in them being the target of numerous market manipulation

attempts. A PD scheme is a particular type of market manipulation strategy that is highly risky for participants [Akyildirim et al., 2022]. Typically, a PD scheme involves the following general steps happening over the course of 2–3 minutes.

- **Pre-Pump:** Organisers discretely advertise to their members which cryptocurrency to buy, which cryptocurrency to use to make the purchases, and the exact time to do so.
- **Pump:** Members are urged to buy the target coin and hold it to cause a short-term price inflation.
- **Dump:** After the price increases dramatically, some members consolidate their gains by selling. This leads to a cascade as more and more participants choose to sell, often leading to the price returning close to its pre-pump level.

A more thorough discussion of PD schemes may be found in subsections 3.3–3.5 in [Akyildirim et al., 2022]. Further details may also be found in [Xu and Livshits, 2019] or [La Morgia et al., 2023], for example.

The authors of [Akyildirim et al., 2022] consider data obtained from the database created by [La Morgia et al., 2023]. The database consists of a list of PD attempts. Every entry of the list contains

- (1) The acronym for the name of the coin, e.g. BTC for Bitcoin, called *symbol* of the coin;
- (2) The name of the Telegram/Discord group where this information was retrieved;
- (3) The date (day) of the PD attempt;
- (4) The time (hour and minute) of the PD attempt;
- (5) The exchange on which the PD attempt took place.

Using this information and the Python-library `ccxt`, the authors of [Akyildirim et al., 2022] download trades, that is all orders that were placed and also realised by selling or buying cryptocurrency in a neighbourhood of the PD attempt. Every trade is characterised by the following information:

- (1) Trading pair of symbol and currency, e.g. BTC/USD means Bitcoin is traded using US Dollars;
- (2) Timestamp, an integer, denoting the UNIX time in milliseconds;
- (3) Datetime, i.e. the ISO8601 datetime with milliseconds;
- (4) Side, string, either “buy” or “sell” to distinguish the operation kind: for the trading pair BTC/USD “buy” means buying BTC using USD, while “sell” means receiving USD for BTC;
- (5) Price, float, indicating the price at which the pair was traded;
- (6) Amount, float, denoting the quantity (in base currency, i.e. USD in the example BTC/USD) which was traded.

Data was retrieved for windows of length 3, 6, and 14 days centred on the moment of the PD attempt. The underlying idea is that the activity corresponding to the PD attempt will be anomalous compared to the normal activity contained within the window.

The raw data was preprocessed as follows. Whenever trades have the same timestamp, side, and price then the trade volumes are summed accumulated since one can interpret these trades as one larger trade without any loss of information. The variable side is translated into a numerical format by encoding all “buy” signals as 0.5 and all “sell” signals as -0.5 . Volume is calculated as the product of price and amount; prices are used to compute the simple returns. Thus the final time series considered consists of the channels trade time, returns, volume, and trade side, with each channel normalised to live within the interval $[0, 1]$.

Let $D_c \subset \mathbb{R}^{n+4}$ denote the vector of time series composed of trades’ timestamp, side, price, and volume for a coin c . Here n denotes the number of trades during the pre-specified time interval around the PD attempt. Split D_c as into $\lceil n/o \rceil$ subsets of w trades, where o denotes the offset of how many trades to move forward in

time. To be more precise, identify the trade-window specified by w trades and then move on shifting among the listed trades. Denote this splitting by

$$D_c = \bigcup_{i=1}^{\lceil n/o \rceil} D_c^i. \quad (15.6)$$

Transform each subset D_c^i into a feature set by either calculating the exact signature or the randomised signature. Let R_c^i and \tilde{R}_c^i denote the exact and randomised signatures respectively. The final feature set is then given by

$$\mathbf{R} := \bigcup_c \bigcup_{i=1}^{\lceil n/o \rceil} R_c^i \quad \text{and} \quad \tilde{\mathbf{R}} := \bigcup_c \bigcup_{i=1}^{\lceil n/o \rceil} \tilde{R}_c^i. \quad (15.7)$$

The offset o and the window size w can be viewed as hyper-parameters which can be optimised. The choice of $w = 100$ and $o = 5$ was found to give a good balance between run time and precision in [Akyildirim et al., 2022].

The remaining hyper-parameters required to compute the randomised signatures were chosen as follows. The function σ is taken as \tanh , the reservoir dimension k was chosen as 50, the mean and the variance of the matrices \mathbf{A} were 0.05 and 0.1 respectively, and the mean and the variance of the vectors b were 0 and 1 respectively (cf. Table 1 in [Akyildirim et al., 2022]).

Once the feature sets \mathbf{R} and $\tilde{\mathbf{R}}$ are computed, the authors of [Akyildirim et al., 2022] use the **Sklearn Python Package** implementations of the robust covariance and isolation forest anomaly detection algorithms. In contrast to the signature-based approach to anomaly detection proposed in [Shao et al., 2020] (which is covered in Section 14 of this article), these “off-the-shelf” anomaly detection algorithms are used *without* alteration in [Akyildirim et al., 2022]. Consequently, any drawbacks of these anomaly detection algorithms are inherited. The method using the robust covariance anomaly detection algorithm requires the number of samples to exceed the square of the number of features, which leads to a restriction on the truncation level that may be considered. The method using the isolation forest anomaly detection algorithm requires the use of sub-sampling techniques due to the improved performance of this algorithm on smaller data sets.

More detailed overviews of the robust covariance and isolation forest anomaly detection algorithms are provided in subsections 3.8.1 and 3.8.2 of [Akyildirim et al., 2022] respectively. The reader seeking extensive coverage is directed to [EStimator, 1999] for the robust covariance anomaly detection algorithm and to [Liu et al., 2008] for the isolation forest anomaly detection algorithm.

The predictions of the resulting methods are compared with manually labelled PD attempts. The unsupervised anomaly detection method proposed in [Kamps and Kleinberg, 2018] is used as a benchmark. Performance is additionally compared against the method proposed in [La Morgia et al., 2023] that is based upon supervised learning and requires labelled training data. The signature and randomised signature methods are themselves unsupervised techniques; no labelled training data is required. The precision, recall, and F1 scores are recorded for comparison (the F1 score is the harmonic mean of the precision and the recall). The results are summarised in Table 19 (cf. Table 3 in [Akyildirim et al., 2022]).

Excluding the Exact Truncated Signature (Robust Covariance) method, all the signature and randomised signature based methods proposed in [Akyildirim et al., 2022] out-perform the benchmark unsupervised model from [Kamps and Kleinberg, 2018] for all time windows. Moreover, the best classifier from [Akyildirim et al., 2022], which is an unsupervised method, comes relatively close to matching the performance of the classifier in [La Morgia et al., 2023] which is a supervised method. Finally, the similar levels of performance between the exact truncated signature methods and the randomised signature methods provides an empirical illustration of the theoretical property that randomised signatures contain the same information as the signature itself.

Classifier	Days	Precision (%)	Recall (%)	F1 (%)
Kamps et al.	3	65	86	74
	6	48	86	62
	14	29	81	43
La Morgia et al.	3	98	91	95
	6	-	-	-
	14	-	-	-
Randomised Signature (Isolation Forest)	3	93	83	88
	6	82	83	82
	14	75	75	75
Randomised Signature (Robust Covariance)	3	80	94	86
	6	80	84	82
	14	75	75	75
Exact Truncated Signature (Isolation Forest)	3	83	92	87
	6	81	82	81
	14	71	74	72
Exact Truncated Signature (Robust Covariance)	3	61	61	60
	6	61	47	54
	14	58	32	44

Table 19: Kamps et al. denotes the method proposed in [Kamps and Kleinberg, 2018]. La Morgia et al. denotes the method proposed in [La Morgia et al., 2023]; the results for this method are reported directly from [La Morgia et al., 2023] which does not consider windows of 6 or 14 days. Randomised Signature (Isolation Forest) denotes the method proposed in [Akyildirim et al., 2022] of using randomised signatures as the features for an isolation forest. Randomised Signature (Robust Covariance) denotes the method proposed in [Akyildirim et al., 2022] of using randomised signatures as the features for the robust covariance algorithm. Exact Truncated Signature (Isolation Forest) denotes the method proposed in [Akyildirim et al., 2022] of using truncated signatures as the features for an isolation forest. Exact Truncated Signature (Robust Covariance) denotes the method proposed in [Akyildirim et al., 2022] of using truncated signatures as the features for the robust covariance algorithm. Table 3 in [Akyildirim et al., 2022].

References

- [Akyildirim et al., 2022] Akyildirim, E., Gambar, M., Teichmann, J., and Zhou, S. (2022). Applications of signature methods to market anomaly detection. *arXiv preprint arXiv:2201.02441*.
- [Aldeneh and Provost, 2017] Aldeneh, Z. and Provost, E. M. (2017). Using regional saliency for speech emotion recognition. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 2741–2745. IEEE.
- [Alpaydin and Alimoglu, 1996] Alpaydin, E. and Alimoglu, F. (1996). Pen-Based Recognition of Handwritten Digits. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5MG6K>.
- [Alsentzer et al., 2019] Alsentzer, E., Murphy, J., Boag, W., Weng, W.-H., Jindi, D., Naumann, T., and McDermott, M. (2019). Publicly available clinical bert embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78.
- [Amer et al., 2013] Amer, M., Goldstein, M., and Abdennadher, S. (2013). 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.
- [Aronszajn, 1950] Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American mathematical society*, 68(3):337–404.
- [Arribas et al., 2020] Arribas, I. P., Salvi, C., and Szpruch, L. (2020). Sig-sdes model for quantitative finance. In *Proceedings of the First ACM International Conference on AI in Finance*, pages 1–8.

- [Arrubarrena et al., 2024] Arrubarrena, P., Lemercier, M., Nikolic, B., Lyons, T., and Cass, T. (2024). Novelty detection on radio astronomy data using signatures. *arXiv preprint arXiv:2402.14892*.
- [Bagnall et al., 2018] Bagnall, A., Dau, H. A., Lines, J., Flynn, M., Large, J., Bostrom, A., Southam, P., and Keogh, E. (2018). The uea multivariate time series classification archive, 2018. *arXiv preprint arXiv:1811.00075*.
- [Bagnall et al., 2017] Bagnall, A., Lines, J., Bostrom, A., Large, J., and Keogh, E. (2017). The great time series classification bake off: A review and experimental evaluation of recent algorithmic advances. *Data mining and knowledge discovery*, 31:606–660.
- [Balvers et al., 2000] Balvers, R., Wu, Y., and Gilliland, E. (2000). Mean reversion across national stock markets and parametric contrarian investment strategies. *The Journal of Finance*, 55(2):745–772.
- [Bank et al., 2023] Bank, P., Bayer, C., Friz, P. K., and Pelizzari, L. (2023). Rough pdes for local stochastic volatility models. *arXiv preprint arXiv:2307.09216*.
- [Bayer et al., 2023a] Bayer, C., Breneis, S., and Lyons, T. (2023a). An adaptive algorithm for rough differential equations. *arXiv preprint arXiv:2307.12590*.
- [Bayer et al., 2023b] Bayer, C., Friz, P. K., and Tapia, N. (2023b). Stability of deep neural networks via discrete rough paths. *SIAM Journal on Mathematics of Data Science*, 5(1):50–76.
- [Bayer et al., 2023c] Bayer, C., Hager, P. P., Riedel, S., and Schoenmakers, J. (2023c). Optimal stopping with signatures. *The Annals of Applied Probability*, 33(1):238–273.
- [Bayer et al., 2023d] Bayer, C., Pelizzari, L., and Schoenmakers, J. (2023d). Primal and dual optimal stopping with signatures. *arXiv preprint arXiv:2312.03444*.
- [Biyong et al., 2020] Biyong, J. P., Wang, B., Lyons, T., and Nevado-Holgado, A. (2020). Information extraction from swedish medical prescriptions with sig-transformer encoder. In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 41–54.
- [Boedihardjo et al., 2021] Boedihardjo, H., Diehl, J., Mezzarobba, M., and Ni, H. (2021). The expected signature of brownian motion stopped on the boundary of a circle has finite radius of convergence. *Bulletin of the London Mathematical Society*, 53(1):285–299.
- [Boedihardjo et al., 2016] Boedihardjo, H., Geng, X., Lyons, T., and Yang, D. (2016). The signature of a rough path: Uniqueness. *Advances in Mathematics*, 293:720–737.
- [Bogachev, 1998] Bogachev, V. I. (1998). *Gaussian Measures*. Number 62 in Mathematical Surveys and Monographs. American Mathematical Soc.
- [Bonnier and Oberhauser, 2020] Bonnier, P. and Oberhauser, H. (2020). Signature cumulants, ordered partitions, and independence of stochastic processes. *Bernoulli*, 26(4):2727–2757.
- [Boutaib et al., 2013] Boutaib, Y., Gyurkó, L. G., Lyons, T., and Yang, D. (2013). Dimension-free euler estimates of rough differential equations. *arXiv preprint arXiv:1307.4708*.
- [Bradbury et al., 2018] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: Composable transformations of Python+NumPy programs.
- [Breunig et al., 2000] Breunig, M. M., Kriegel, H.-P., Ng, R. T., and Sander, J. (2000). Lof: Identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pages 93–104.
- [Buchman et al., 2020a] Buchman, T. G., Simpson, S. Q., Sciarretta, K. L., Finne, K. P., Sowers, N., Collier, M., Chavan, S., Oke, I., Pennini, M. E., Santhosh, A., et al. (2020a). Sepsis among medicare beneficiaries: 1. the burdens of sepsis, 2012–2018. *Critical care medicine*, 48(3):276–288.

- [Buchman et al., 2020b] Buchman, T. G., Simpson, S. Q., Sciarretta, K. L., Finne, K. P., Sowers, N., Collier, M., Chavan, S., Oke, I., Pennini, M. E., Santhosh, A., et al. (2020b). Sepsis among medicare beneficiaries: 2. the trajectories of sepsis, 2012–2018. *Critical care medicine*, 48(3):289–301.
- [Buchman et al., 2020c] Buchman, T. G., Simpson, S. Q., Sciarretta, K. L., Finne, K. P., Sowers, N., Collier, M., Chavan, S., Oke, I., Pennini, M. E., Santhosh, A., et al. (2020c). Sepsis among medicare beneficiaries: 3. the methods, models, and forecasts of sepsis, 2012–2018. *Critical care medicine*, 48(3):302–318.
- [Busso et al., 2008] Busso, C., Bulut, M., Lee, C.-C., Kazemzadeh, A., Mower, E., Kim, S., Chang, J. N., Lee, S., and Narayanan, S. S. (2008). Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359.
- [Cass et al., 2024] Cass, T., Gonon, L., and Zozoulenko, N. (2024). Variance norms for kernelized anomaly detection. *arXiv preprint arXiv:2407.11873*.
- [Cass et al., 2021] Cass, T., Lyons, T., and Xu, X. (2021). General signature kernels. *arXiv preprint arXiv:2107.00447*.
- [Cass and Turner, 2024] Cass, T. and Turner, W. F. (2024). Topologies on unparameterised path space. *Journal of Functional Analysis*, 286(4):110261.
- [Chalapathy and Chawla, 2019] Chalapathy, R. and Chawla, S. (2019). Deep learning for anomaly detection: A survey. *arXiv preprint arXiv:1901.03407*.
- [Chen et al., 2018] Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary differential equations. *Advances in neural information processing systems*, 31.
- [Chéron et al., 2015] Chéron, G., Laptev, I., and Schmid, C. (2015). P-cnn: Pose-based cnn features for action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3218–3226.
- [Chevyrev, 2014] Chevyrev, I. (2014). Unitary representations of geometric rough paths. *arXiv preprint arXiv:1307.3580*.
- [Chevyrev and Kormilitzin, 2016] Chevyrev, I. and Kormilitzin, A. (2016). A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*.
- [Chevyrev and Lyons, 2016] Chevyrev, I. and Lyons, T. (2016). Characteristic functions of measures on geometric rough paths. *The Annals of Probability*, 44(6):4049–4082.
- [Chevyrev and Oberhauser, 2022] Chevyrev, I. and Oberhauser, H. (2022). Signature moments to characterize laws of stochastic processes. *Journal of Machine Learning Research*, 23(176):1–42.
- [Cohen et al., 2023] Cohen, S. N., Lui, S., Malpass, W., Mantoan, G., Nesheim, L., de Paula, A., Reeves, A., Scott, C., Small, E., and Yang, L. (2023). Nowcasting with signature methods. *arXiv preprint arXiv:2305.10256*.
- [Cuchiero et al., 2023a] Cuchiero, C., Gazzani, G., and Svaluto-Ferro, S. (2023a). Signature-based models: Theory and calibration. *SIAM journal on financial mathematics*, 14(3):910–957.
- [Cuchiero et al., 2021a] Cuchiero, C., Gonon, L., Grigoryeva, L., Ortega, J.-P., and Teichmann, J. (2021a). Discrete-time signatures and randomness in reservoir computing. *IEEE Transactions on Neural Networks and Learning Systems*, 33(11):6321–6330.
- [Cuchiero et al., 2021b] Cuchiero, C., Gonon, L., Grigoryeva, L., Ortega, J.-P., and Teichmann, J. (2021b). Expressive power of randomized signature. In *The Symbiosis of Deep Learning and Differential Equations*.
- [Cuchiero and Möller, 2023] Cuchiero, C. and Möller, J. (2023). Signature methods in stochastic portfolio theory. *arXiv preprint arXiv:2310.02322*.

- [Cuchiero et al., 2022] Cuchiero, C., Primavera, F., and Svaluto-Ferro, S. (2022). Universal approximation theorems for continuous functions of càdlàg paths and Lévy-type signature models. *arXiv preprint arXiv:2208.02293*.
- [Cuchiero et al., 2023b] Cuchiero, C., Schmock, P., and Teichmann, J. (2023b). Global universal approximation of functional input maps on weighted spaces. *arXiv preprint arXiv:2306.03303*.
- [Cuchiero et al., 2023c] Cuchiero, C., Svaluto-Ferro, S., and Teichmann, J. (2023c). Signature sdes from an affine and polynomial perspective. *arXiv preprint arXiv:2302.01362*.
- [Cuturi, 2011] Cuturi, M. (2011). Fast global alignment kernels. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 929–936.
- [Cuturi et al., 2007] Cuturi, M., Vert, J.-P., Birkenes, O., and Matsui, T. (2007). A kernel for time series based on global alignments. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 2, pages II–413. IEEE.
- [Dahikar and Rode, 2014] Dahikar, S. S. and Rode, S. V. (2014). 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.
- [Du et al., 2015] Du, Y., Wang, W., and Wang, L. (2015). Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118.
- [Dupire and Tissot-Daguette, 2022] Dupire, B. and Tissot-Daguette, V. (2022). Functional expansions. *arXiv preprint arXiv:2212.13628*.
- [Dupire and Tissot-Daguette, 2025] Dupire, B. and Tissot-Daguette, V. (2025). Signature and the functional Taylor expansion. In *(Signature methods in Finance)*. Springer (to appear).
- [EStimator, 1999] EStimator, D. (1999). A fast algorithm for the minimum covariance. *Technometrics*, 41(3):212.
- [Eyben et al., 2015] Eyben, F., Scherer, K. R., Schuller, B. W., Sundberg, J., André, E., Busso, C., Devillers, L. Y., Epps, J., Laukka, P., Narayanan, S. S., et al. (2015). The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing. *IEEE transactions on affective computing*, 7(2):190–202.
- [Eyben et al., 2010] Eyben, F., Wöllmer, M., and Schuller, B. (2010). Opensmile: The munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1459–1462.
- [Fan et al., 2016] Fan, J., Zha, Z., and Tian, X. (2016). Action recognition with novel high-level pose features. In *2016 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6. IEEE.
- [Fang et al., 2017] Fang, H.-S., Xie, S., Tai, Y.-W., and Lu, C. (2017). Rmpe: Regional multi-person pose estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 2334–2343.
- [Fawcett, 2003] Fawcett, T. (2003). *Problems in Stochastic Analysis. Connections Between Rough Paths and Non-Commutative Harmonic Analysis*. PhD thesis, University of Oxford.
- [Flaxman, 2015] Flaxman, S. R. (2015). *Machine Learning in Space and Time*. PhD thesis, Ph. D. thesis, Carnegie Mellon University.
- [Foster et al., 2020] Foster, J., Lyons, T., and Oberhauser, H. (2020). An optimal polynomial approximation of brownian motion. *SIAM Journal on Numerical Analysis*, 58(3):1393–1421.
- [Freedman, 2012] Freedman, D. (2012). *Brownian Motion and Diffusion*. Springer Science & Business Media.

- [Friz and Victoir, 2006] Friz, P. and Victoir, N. (2006). A note on the notion of geometric rough paths. *Probability theory and related fields*, 136(3):395–416.
- [Friz, 2005] Friz, P. K. (2005). Continuity of the itô-map for hölder rough paths with applications to the support theorem in hölder norm. *Probability and partial differential equations in modern applied mathematics*, pages 117–135.
- [Friz et al., 2022] Friz, P. K., Hager, P. P., and Tapia, N. (2022). Unified signature cumulants and generalized magnus expansions. In *Forum of Mathematics, Sigma*, volume 10, page e42. Cambridge University Press.
- [Friz and Victoir, 2010] Friz, P. K. and Victoir, N. B. (2010). *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*, volume 120. Cambridge University Press.
- [Furey et al., 2000] Furey, T. S., Cristianini, N., Duffy, N., Bednarski, D. W., Schummer, M., and Haussler, D. (2000). Support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics*, 16(10):906–914.
- [Gatheral et al., 2018] Gatheral, J., Jaisson, T., and Rosenbaum, M. (2018). Volatility is rough. *Quantitative finance*, 18(6):933–949.
- [Géron, 2022] Géron, A. (2022). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*. ” O’Reilly Media, Inc.”.
- [Ghosh et al., 2016] Ghosh, S., Laksana, E., Morency, L.-P., and Scherer, S. (2016). Representation learning for speech emotion recognition. In *Interspeech*, pages 3603–3607.
- [Grigoryeva and Ortega, 2018] Grigoryeva, L. and Ortega, J.-P. (2018). Echo state networks are universal. *Neural Networks*, 108:495–508.
- [Gu et al., 2021] Gu, Y., Tinn, R., Cheng, H., Lucas, M., Usuyama, N., Liu, X., Naumann, T., Gao, J., and Poon, H. (2021). Domain-specific language model pretraining for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3(1):1–23.
- [Gurney, 1997] Gurney, K. (1997). *An Introduction to Neural Networks*. CRC Press.
- [Hambly and Lyons, 2010] Hambly, B. and Lyons, T. (2010). Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, 171(1):109–167.
- [Han et al., 2018] Han, W., Ruan, H., Chen, X., Wang, Z., Li, H., and Schuller, B. (2018). Towards temporal modelling of categorical speech emotion recognition. *Interspeech 2018*.
- [Hastie et al., 2017] Hastie, T., Tibshirani, R., and Friedman, J. (2017). The elements of statistical learning: Data mining, inference, and prediction.
- [Hauskrecht et al., 2013] Hauskrecht, M., Batal, I., Valko, M., Visweswaran, S., Cooper, G. F., and Clermont, G. (2013). Outlier detection for patient monitoring and alerting. *Journal of biomedical informatics*, 46(1):47–55.
- [Hautamaki et al., 2004] Hautamaki, V., Karkkainen, I., and Franti, P. (2004). Outlier detection using k-nearest neighbour graph. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 3, pages 430–433. IEEE.
- [He and Wang, 2007] He, Q. P. and Wang, J. (2007). Fault detection using the k-nearest neighbor rule for semiconductor manufacturing processes. *IEEE Transactions on Semiconductor Manufacturing*, 20(4):345–354.
- [He et al., 2003] He, Z., Xu, X., and Deng, S. (2003). Discovering cluster-based local outliers. *Pattern recognition letters*, 24(9-10):1641–1650.
- [Heaton, 2011] Heaton, J. (2011). Introduction to the math of neural networks (beta-1). *Heaton Research Inc.*

- [Hill, 1986] Hill, T. L. (1986). *An Introduction to Statistical Thermodynamics*. Courier Corporation.
- [Hodgkinson et al., 2020] Hodgkinson, L., van der Heide, C., Roosta, F., and Mahoney, M. W. (2020). Stochastic normalizing flows. *arXiv preprint arXiv:2002.09547*.
- [Horvath et al., 2023] Horvath, B., Lemercier, M., Liu, C., Lyons, T., and Salvi, C. (2023). Optimal stopping via distribution regression: A higher rank signature approach. *arXiv preprint arXiv:2304.01479*.
- [Hu et al., 2015] Hu, J.-F., Zheng, W.-S., Lai, J., and Zhang, J. (2015). Jointly learning heterogeneous features for rgb-d activity recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5344–5352.
- [Huang et al., 2019] Huang, K., Altosaar, J., and Ranganath, R. (2019). Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*.
- [Hurst, 1951] Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American society of civil engineers*, 116(1):770–799.
- [Jhuang et al., 2013] Jhuang, H., Gall, J., Zuffi, S., Schmid, C., and Black, M. J. (2013). Towards understanding action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3192–3199.
- [Ji et al., 2014] Ji, Y., Ye, G., and Cheng, H. (2014). Interactive body part contrast mining for human interaction recognition. In *2014 IEEE international conference on multimedia and expo workshops (ICMEW)*, pages 1–6. IEEE.
- [Johansson, 1973] Johansson, G. (1973). Visual perception of biological motion and a model for its analysis. *Perception & psychophysics*, 14:201–211.
- [Jones and Sielken, 2000] Jones, A. K. and Sielken, R. S. (2000). Computer system intrusion detection: A survey. *Computer Science Technical Report*, pages 1–25.
- [Kamps and Kleinberg, 2018] Kamps, J. and Kleinberg, B. (2018). To the moon: Defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1):1–18.
- [Kapsouras and Nikolaidis, 2014] Kapsouras, I. and Nikolaidis, N. (2014). Action recognition on motion capture data using a dynemes and forward differences representation. *Journal of Visual Communication and Image Representation*, 25(6):1432–1445.
- [Ke et al., 2017] Ke, Q., An, S., Bennamoun, M., Sohel, F., and Boussaid, F. (2017). Skeletonnet: Mining deep part features for 3-d action recognition. *IEEE signal processing letters*, 24(6):731–735.
- [Kenton and Toutanova, 2019] Kenton, J. D. M.-W. C. and Toutanova, L. K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019*, number 2 in *Proceedings of NAACL-HLT*. Minneapolis, Minnesota, Association for Computational Linguistics.
- [Kidger et al., 2019] Kidger, P., Bonnier, P., Perez Arribas, I., Salvi, C., and Lyons, T. (2019). Deep signature transforms. *Advances in Neural Information Processing Systems*, 32.
- [Kidger et al., 2021a] Kidger, P., Foster, J., Li, X., and Lyons, T. J. (2021a). Neural sdes as infinite-dimensional gans. In *International conference on machine learning*, pages 5453–5463. PMLR.
- [Kidger et al., 2021b] Kidger, P., Foster, J., Li, X. C., and Lyons, T. (2021b). Efficient and accurate gradients for neural sdes. *Advances in Neural Information Processing Systems*, 34:18747–18761.
- [Kidger and Lyons, 2020] Kidger, P. and Lyons, T. (2020). Signatory: Differentiable computations of the signature and logsignature transforms, on both cpu and gpu. In *International Conference on Learning Representations*.

- [Kidger et al., 2020] Kidger, P., Morrill, J., Foster, J., and Lyons, T. (2020). Neural controlled differential equations for irregular time series. *Advances in Neural Information Processing Systems*, 33:6696–6707.
- [Király and Oberhauser, 2019] Király, F. J. and Oberhauser, H. (2019). Kernels for sequentially ordered data. *Journal of Machine Learning Research*, 20(31):1–45.
- [Kreicbergs and Grislis, 2019] Kreicbergs, J. and Grislis, A. (2019). Periodic polynomial regression analysis of urban driving characteristics. *Argonomy Research*, 17:1080–1088.
- [Kuelbs, 1970] Kuelbs, J. (1970). Gaussian measures on a banach space. *Journal of Functional Analysis*, 5(3):354–367.
- [Kumar et al., 2006] Kumar, A., Roberts, D., Wood, K. E., Light, B., Parrillo, J. E., Sharma, S., Suppes, R., Feinstein, D., Zanotti, S., Taiberg, L., et al. (2006). Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock. *Critical care medicine*, 34(6):1589–1596.
- [La Morgia et al., 2023] La Morgia, M., Mei, A., Sassi, F., and Stefa, J. (2023). The doge of wall street: Analysis and detection of pump and dump cryptocurrency manipulations. *ACM Transactions on Internet Technology*, 23(1):1–28.
- [Lacasa et al., 2009] Lacasa, L., Luque, B., Luque, J., and Nuno, J. C. (2009). The visibility graph: A new method for estimating the hurst exponent of fractional brownian motion. *Europhysics Letters*, 86(3):30001.
- [Law et al., 2018] Law, H. C. L., Sutherland, D. J., Sejdinovic, D., and Flaxman, S. (2018). Bayesian approaches to distribution regression. In *International Conference on Artificial Intelligence and Statistics*, pages 1167–1176. PMLR.
- [Lee and Tashev, 2015] Lee, J. and Tashev, I. (2015). High-level feature representation using recurrent neural network for speech emotion recognition. In *Interspeech 2015*.
- [Lees, 1960] Lees, M. (1960). The goursat problem. *Journal of the Society for Industrial and Applied Mathematics*, 8(3):518–530.
- [Lemerrier et al., 2021] Lemerrier, M., Salvi, C., Damoulas, T., Bonilla, E., and Lyons, T. (2021). Distribution regression for sequential data. In *International Conference on Artificial Intelligence and Statistics*, pages 3754–3762. PMLR.
- [Levin et al., 2013] Levin, D., Lyons, T., and Ni, H. (2013). Learning from the past, predicting the statistics for the future, learning an evolving system. *arXiv preprint arXiv:1309.0260*.
- [Li et al., 2017] Li, C., Hou, Y., Wang, P., and Li, W. (2017). Joint distance maps based action recognition with convolutional neural networks. *IEEE Signal Processing Letters*, 24(5):624–628.
- [Li et al., 2015] Li, W., Wen, L., Chuah, M. C., and Lyu, S. (2015). Category-blind human action recognition: A practical recognition system. In *Proceedings of the IEEE international conference on computer vision*, pages 4444–4452.
- [Li et al., 2020] Li, X., Wong, T.-K. L., Chen, R. T., and Duvenaud, D. K. (2020). Scalable gradients and variational inference for stochastic differential equations. In *Symposium on Advances in Approximate Bayesian Inference*, pages 1–28. PMLR.
- [Lifshits and Lifshits, 2012] Lifshits, M. and Lifshits, M. (2012). *Lectures on Gaussian Processes*. Springer.
- [Litterer and Lyons, 2007] Litterer, C. and Lyons, T. (2007). Cubature on Wiener Space Continued. In Aka-hori, J., Ogawa, S., and Watanabe, S., editors, *Stochastic Processes And Applications To Mathematical Finance*, World Scientific Book Chapters, chapter 12, pages 197–217. World Scientific Publishing Co. Pte. Ltd.

- [Liu et al., 2008] Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2008). Isolation forest. In *2008 eighth IEEE international conference on data mining*, pages 413–422. IEEE.
- [Liu et al., 2012] Liu, F. T., Ting, K. M., and Zhou, Z.-H. (2012). Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 6(1):1–39.
- [Liu et al., 2017a] Liu, J., Shahroudy, A., Xu, D., Kot, A. C., and Wang, G. (2017a). Skeleton-based action recognition using spatio-temporal lstm network with trust gates. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):3007–3021.
- [Liu et al., 2016] Liu, J., Shahroudy, A., Xu, D., and Wang, G. (2016). Spatio-temporal lstm with trust gates for 3d human action recognition. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 816–833. Springer.
- [Liu et al., 2017b] Liu, J., Wang, G., Hu, P., Duan, L.-Y., and Kot, A. C. (2017b). Global context-aware attention lstm networks for 3d action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1647–1656.
- [Loranger et al., 1994] Loranger, A. W., Sartorius, N., Andreoli, A., Berger, P., Buchheim, P., Channabasavanna, S., Coid, B., Dahl, A., Diekstra, R. F., Ferguson, B., et al. (1994). 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*, 51(3):215–224.
- [Lyons,] Lyons, T. Esig python package.
- [Lyons, 2014] Lyons, T. (2014). Rough paths, signatures and the modelling of functions on streams. *arXiv preprint arXiv:1405.4537*.
- [Lyons and Ni, 2015] Lyons, T. and Ni, H. (2015). Expected signature of brownian motion up to the first exit time from a bounded domain. *The Annals of Probability*, 43(5):2729 – 2762.
- [Lyons and Qian, 2002] Lyons, T. and Qian, Z. (2002). *System Control and Rough Paths*. Oxford University Press.
- [Lyons and Victoir, 2004] Lyons, T. and Victoir, N. (2004). Cubature on wiener space. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 460(2041):169–198.
- [Lyons, 1998] Lyons, T. J. (1998). Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310.
- [Lyons et al., 2007] Lyons, T. J., Caruana, M., and Lévy, T. (2007). *Differential Equations Driven By Rough Paths*. Springer.
- [Lyons and Yang, 2013] Lyons, T. J. and Yang, D. (2013). 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(12):3067–3103.
- [Mahalanobis, 2018] Mahalanobis, P. C. (2018). On the generalized distance in statistics. *Sankhyā: The Indian Journal of Statistics, Series A (2008-)*, 80:S1–S7.
- [Manton et al., 2015] Manton, J. H., Amblard, P.-O., et al. (2015). A primer on reproducing kernel hilbert spaces. *Foundations and Trends® in Signal Processing*, 8(1–2):1–126.
- [Mirsamadi et al., 2017] Mirsamadi, S., Barsoum, E., and Zhang, C. (2017). Automatic speech emotion recognition using recurrent neural networks with local attention. In *2017 IEEE International conference on acoustics, speech and signal processing (ICASSP)*, pages 2227–2231. IEEE.
- [Moore et al., 2019] Moore, P., Lyons, T., Gallacher, J., and Initiative, A. D. N. (2019). Using path signatures to predict a diagnosis of alzheimer’s disease. *PloS one*, 14(9):e0222212.

- [Morley and Lyons, 2024] Morley, S. and Lyons, T. (2024). Roughpy: Streaming data is rarely smooth. *Proceedings of the 23rd Python in Science Conference*.
- [Morrill et al., 2021a] Morrill, J., Kidger, P., Yang, L., and Lyons, T. (2021a). Neural controlled differential equations for online prediction tasks. *arXiv preprint arXiv:2106.11028*.
- [Morrill et al., 2019] Morrill, J., Kormilitzin, A., Nevado-Holgado, A., Swaminathan, S., Howison, S., and Lyons, T. (2019). The signature-based model for early detection of sepsis from electronic health records in the intensive care unit. In *2019 Computing in Cardiology (CinC)*, pages Page–1. IEEE.
- [Morrill et al., 2021b] Morrill, J., Salvi, C., Kidger, P., and Foster, J. (2021b). Neural rough differential equations for long time series. In *International Conference on Machine Learning*, pages 7829–7838. PMLR.
- [Morrill et al., 2020] Morrill, J. H., Kormilitzin, A., Nevado-Holgado, A. J., Swaminathan, S., Howison, S. D., and Lyons, T. J. (2020). 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):e976–e981.
- [Muandet et al., 2012] Muandet, K., Fukumizu, K., Dinuzzo, F., and Schölkopf, B. (2012). Learning from distributions via support measure machines. *Advances in neural information processing systems*, 25.
- [Neufeld and Schmocker, 2022] Neufeld, A. and Schmocker, P. (2022). Chaotic hedging with iterated integrals and neural networks. *arXiv preprint arXiv:2209.10166*.
- [Neumann and Vu, 2017] Neumann, M. and Vu, N. T. (2017). Attentive convolutional neural network based speech emotion recognition: A study on the impact of input features, signal length, and acted speech. *arXiv preprint arXiv:1706.00612*.
- [Nian et al., 2016] Nian, K., Zhang, H., Tayal, A., Coleman, T., and Li, Y. (2016). Auto insurance fraud detection using unsupervised spectral ranking for anomaly. *The Journal of Finance and Data Science*, 2(1):58–75.
- [Offringa et al., 2012] Offringa, A., Van De Gronde, J., and Roerdink, J. (2012). A morphological algorithm for improving radio-frequency interference detection. *Astronomy & astrophysics*, 539:A95.
- [Offringa et al., 2015] Offringa, A., Wayth, R., Hurley-Walker, N., Kaplan, D., Barry, N., Beardsley, A., Bell, M., Bernardi, G., Bowman, J., Briggs, F., et al. (2015). The low-frequency environment of the murchison widefield array: Radio-frequency interference analysis and mitigation. *Publications of the Astronomical Society of Australia*, 32:e008.
- [Ofli et al., 2013] Ofli, F., Chaudhry, R., Kurillo, G., Vidal, R., and Bajcsy, R. (2013). Berkeley mhad: A comprehensive multimodal human action database. In *2013 IEEE workshop on applications of computer vision (WACV)*, pages 53–60. IEEE.
- [Panda et al., 2010] Panda, S. S., Ames, D. P., and Panigrahi, S. (2010). Application of vegetation indices for agricultural crop yield prediction using neural network techniques. *Remote sensing*, 2(3):673–696.
- [Papavasiliou and Ladrone, 2011] Papavasiliou, A. and Ladrone, C. (2011). Parameter estimation for rough differential equations. *The Annals of Statistics*, 39(4):2047–2073.
- [Paszke et al., 2019] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- [Paulsen and Raghupathi, 2016] Paulsen, V. I. and Raghupathi, M. (2016). *An Introduction to the Theory of Reproducing Kernel Hilbert Spaces*, volume 152. Cambridge university press.

- [Peng et al., 2019] Peng, Y., Yan, S., and Lu, Z. (2019). Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets. *BioNLP 2019*, page 58.
- [Perez Arribas et al., 2020] Perez Arribas, I., Cochrane, T., Foster, P., and Lyons, T. (2020). Anomaly detection on streamed data. *arXiv preprint arXiv:2006.03487v1*.
- [Perez Arribas et al., 2018] Perez Arribas, I., Goodwin, G. M., Geddes, J. R., Lyons, T., and Saunders, K. E. (2018). A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder. *Translational psychiatry*, 8(1):274.
- [Peters et al., 2018] Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In Walker, M., Ji, H., and Stent, A., editors, *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- [Pimentel et al., 2014] Pimentel, M. A., Clifton, D. A., Clifton, L., and Tarassenko, L. (2014). A review of novelty detection. *Signal processing*, 99:215–249.
- [Raschka et al., 2022] Raschka, S., Liu, Y. H., and Mirjalili, V. (2022). *Machine Learning with PyTorch and Scikit-Learn: Develop Machine Learning and Deep Learning Models with Python*. Packt Publishing Ltd.
- [Reichl, 2016] Reichl, L. E. (2016). *A Modern Course in Statistical Physics*. John Wiley & Sons.
- [Reizenstein and Graham, 2018] Reizenstein, J. and Graham, B. (2018). The iisignature library: Efficient calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1802.08252*.
- [Reyna et al., 2020] Reyna, M. A., Josef, C. S., Jeter, R., Shashikumar, S. P., Westover, M. B., Nemati, S., Clifford, G. D., and Sharma, A. (2020). Early prediction of sepsis from clinical data: the physionet/computing in cardiology challenge 2019. *Critical care medicine*, 48(2):210–217.
- [Rubanova et al., 2019] Rubanova, Y., Chen, R. T., and Duvenaud, D. K. (2019). Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*, 32.
- [Ryan, 2002] Ryan, R. (2002). Introduction to tensor products of banach spaces. *Springer Monographs in Mathematics/Springer-Verlag London, Ltd.*
- [Salvi et al., 2021] Salvi, C., Cass, T., Foster, J., Lyons, T., and Yang, W. (2021). The signature kernel is the solution of a goursat pde. *SIAM Journal on Mathematics of Data Science*, 3(3):873–899.
- [Sapankevych and Sankar, 2009] Sapankevych, N. I. and Sankar, R. (2009). Time series prediction using support vector machines: A survey. *IEEE computational intelligence magazine*, 4(2):24–38.
- [Sarmadi and Karamodin, 2020] Sarmadi, H. and Karamodin, A. (2020). 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*, 140:106495.
- [Schatten, 1985] Schatten, R. (1985). *A Theory of Cross-Spaces*. Number 26 in Annals of Mathematics Studies. Princeton university press.
- [Schlegel, 2021] Schlegel, K. (2021). When is there a representer theorem? reflexive banach spaces. *Advances in Computational Mathematics*, 47(4):54.
- [Schmitter et al., 2015] Schmitter, D., Roche, A., Maréchal, B., Ribes, D., Abdulkadir, A., Bach-Cuadra, M., Daducci, A., Granziera, C., Klöppel, S., Maeder, P., et al. (2015). An evaluation of volume-based morphometry for prediction of mild cognitive impairment and alzheimer’s disease. *NeuroImage: Clinical*, 7:7–17.
- [Schölkopf et al., 2001] Schölkopf, B., Herbrich, R., and Smola, A. J. (2001). A generalized representer theorem. In *International conference on computational learning theory*, pages 416–426. Springer.

- [Scholkopf and Smola, 2018] Scholkopf, B. and Smola, A. J. (2018). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press.
- [Schrödinger, 1989] Schrödinger, E. (1989). *Statistical Thermodynamics*. Courier Corporation.
- [Shahroudy et al., 2016] Shahroudy, A., Liu, J., Ng, T.-T., and Wang, G. (2016). Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019.
- [Shao et al., 2020] Shao, Z., Chan, R. S.-Y., Cochrane, T., Foster, P., and Lyons, T. (2020). Dimensionless anomaly detection on multivariate streams with variance norm and path signature. *arXiv preprint arXiv:2006.03487*.
- [Shmelev and Salvi, 2024] Shmelev, D. and Salvi, C. (2024). Sparse signature coefficient recovery via kernels. *arXiv preprint arXiv:2412.08579*.
- [Shotton et al., 2013] Shotton, J., Sharp, T., Kipman, A., Fitzgibbon, A., Finocchio, M., Blake, A., Cook, M., and Moore, R. (2013). Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124.
- [Si et al., 2020] Si, Y., Chen, Z., Sun, J., Zhang, D., and Qian, P. (2020). A data-driven fault detection framework using mahalanobis distance based dynamic time warping. *IEEE Access*, 8:108359–108370.
- [Si et al., 2019] Si, Y., Wang, J., Xu, H., and Roberts, K. (2019). Enhancing clinical concept extraction with contextual embeddings. *Journal of the American Medical Informatics Association*, 26(11):1297–1304.
- [Singer et al., 2016] Singer, M., Deutschman, C. S., Seymour, C. W., Shankar-Hari, M., Annane, D., Bauer, M., Bellomo, R., Bernard, G. R., Chiche, J.-D., Coopersmith, C. M., et al. (2016). The third international consensus definitions for sepsis and septic shock (sepsis-3). *Jama*, 315(8):801–810.
- [Singh and Bawa, 2022] Singh, H. and Bawa, S. (2022). Predicting covid-19 statistics using machine learning regression model: Li-muli-poly. *Multimedia Systems*, 28(1):113–120.
- [Sipilainen, 1993] Sipilainen, E.-M. (1993). *Pathwise View on Solutions of Stochastic Differential Equations*. PhD thesis, Ph. D. thesis, The University of Edinburgh.
- [Smola et al., 2007] Smola, A., Gretton, A., Song, L., and Schölkopf, B. (2007). A hilbert space embedding for distributions. In *International conference on algorithmic learning theory*, pages 13–31. Springer.
- [Song et al., 2017] Song, S., Lan, C., Xing, J., Zeng, W., and Liu, J. (2017). An end-to-end spatio-temporal attention model for human action recognition from skeleton data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- [Sørensen et al., 2017] Sørensen, L., Igel, C., Pai, A., Balas, I., Anker, C., Lillholm, M., Nielsen, M., Initiative, A. D. N., et al. (2017). Differential diagnosis of mild cognitive impairment and alzheimer’s disease using structural mri cortical thickness, hippocampal shape, hippocampal texture, and volumetry. *NeuroImage: Clinical*, 13:470–482.
- [Stone, 1948] Stone, M. H. (1948). The generalized weierstrass approximation theorem. *Mathematics Magazine*, 21(5):237–254.
- [Stroock and Varadhan, 1997] Stroock, D. W. and Varadhan, S. S. (1997). *Multidimensional Diffusion Processes*, volume 233. Springer Science & Business Media.
- [Szabó et al., 2016] Szabó, Z., Sriperumbudur, B. K., Póczos, B., and Gretton, A. (2016). Learning theory for distribution regression. *Journal of Machine Learning Research*, 17(152):1–40.
- [Tao and Liu, 2018] Tao, F. and Liu, G. (2018). Advanced lstm: A study about better time dependency modeling in emotion recognition. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2906–2910. IEEE.

- [Tong,] Tong, A. Signax: Computing signatures in jax.
- [Tong and Chang, 2001] Tong, S. and Chang, E. (2001). Support vector machine active learning for image retrieval. In *Proceedings of the ninth ACM international conference on Multimedia*, pages 107–118.
- [Tong and Koller, 2001] Tong, S. and Koller, D. (2001). Support vector machine active learning with applications to text classification. *Journal of machine learning research*, 2(Nov):45–66.
- [Toth and Oberhauser, 2020] Toth, C. and Oberhauser, H. (2020). Bayesian learning from sequential data using gaussian processes with signature covariances. In *International Conference on Machine Learning*, pages 9548–9560. PMLR.
- [Tsanas et al., 2016] Tsanas, A., Saunders, K., Bilderbeck, A., Palmius, N., Osipov, M., Clifford, G., Goodwin, G., and De Vos, M. (2016). Daily longitudinal self-monitoring of mood variability in bipolar disorder and borderline personality disorder. *Journal of affective disorders*, 205:225–233.
- [Tseriotou et al., 2023a] Tseriotou, T., Chan, R. S.-Y., Tsakalidis, A., Bilal, I. M., Kochkina, E., Lyons, T., and Liakata, M. (2023a). Sig-networks toolkit: Signature networks for longitudinal language modelling. *arXiv preprint arXiv:2312.03523*.
- [Tseriotou et al., 2023b] Tseriotou, T., Tsakalidis, A., Foster, P., Lyons, T., and Liakata, M. (2023b). Sequential path signature networks for personalised longitudinal language modeling. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5016–5031.
- [Tzen and Raginsky, 2019] Tzen, B. and Raginsky, M. (2019). Neural stochastic differential equations: Deep latent gaussian models in the diffusion limit. *arXiv preprint arXiv:1905.09883*.
- [Vantigodi and Babu, 2013] Vantigodi, S. and Babu, R. V. (2013). Real-time human action recognition from motion capture data. In *2013 Fourth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG)*, pages 1–4. IEEE.
- [Vantigodi and Radhakrishnan, 2014] Vantigodi, S. and Radhakrishnan, V. B. (2014). Action recognition from motion capture data using meta-cognitive rbf network classifier. In *2014 IEEE Ninth International Conference on Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP)*, pages 1–6. IEEE.
- [Vapnik, 1998] Vapnik, V. (1998). The support vector method of function estimation. In *Nonlinear modeling: Advanced black-box techniques*, pages 55–85. Springer.
- [Varzaneh et al., 2022] Varzaneh, M. G., Riedel, S., Schmeding, A., and Tapia, N. (2022). The geometry of controlled rough paths. *arXiv preprint arXiv:2203.05946*.
- [Vemulapalli et al., 2014] Vemulapalli, R., Arrate, F., and Chellappa, R. (2014). Human action recognition by representing 3d skeletons as points in a lie group. *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 588–595.
- [Walker et al., 2024] Walker, B., Mcleod, A. D., Qin, T., Cheng, Y., Li, H., and Lyons, T. (2024). Log neural controlled differential equations: The lie brackets make a difference. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F., editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 49822–49844. PMLR.
- [Wan et al., 2013] Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., and Fergus, R. (2013). Regularization of neural networks using dropconnect. In *International conference on machine learning*, pages 1058–1066. PMLR.
- [Wang et al., 2019] Wang, B., Liakata, M., Ni, H., Lyons, T., Nevado-Holgado, A. J., and Saunders, K. (2019). A path signature approach for speech emotion recognition. In *Interspeech 2019*, pages 1661–1665. ISCA.

- [Wang et al., 2020] Wang, B., Wu, Y., Taylor, N., Lyons, T., Liakata, M., Nevado-Holgado, A., and Saunders, K. (2020). Learning to detect bipolar disorder and borderline personality disorder with language and speech in non-clinical interviews. In *Proceedings of Interspeech 2020*. International Speech Communication Association.
- [Wang et al., 2021] Wang, B., Wu, Y., Vaci, N., Liakata, M., Lyons, T., and Saunders, K. E. (2021). Modelling paralinguistic properties in conversational speech to detect bipolar disorder and borderline personality disorder. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7243–7247. IEEE.
- [Wang et al., 2013] Wang, C., Wang, Y., and Yuille, A. L. (2013). An approach to pose-based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 915–922.
- [Wang and Schmid, 2013] Wang, H. and Schmid, C. (2013). Action recognition with improved trajectories. In *Proceedings of the IEEE international conference on computer vision*, pages 3551–3558.
- [Westman et al., 2011] Westman, E., Simmons, A., Muehlboeck, J.-S., Mecocci, P., Vellas, B., Tsolaki, M., Kłoszewska, I., Soininen, H., Weiner, M. W., Lovestone, S., et al. (2011). Addneuromed and adni: Similar patterns of alzheimer’s atrophy and automated mri classification accuracy in europe and north america. *Neuroimage*, 58(3):818–828.
- [Wilensky et al., 2019] Wilensky, M. J., Morales, M. F., Hazelton, B. J., Barry, N., Byrne, R., and Roy, S. (2019). Absolving the sins of precision interferometric radio data: A new technique for mitigating faint radio frequency interference. *Publications of the Astronomical Society of the Pacific*, 131(1005):114507.
- [Wu et al., 2022] Wu, Y., Goodwin, G. M., Lyons, T., and Saunders, K. E. (2022). Identifying psychiatric diagnosis from missing mood data through the use of log-signature features. *Plos one*, 17(11):e0276821.
- [Xu and Livshits, 2019] Xu, J. and Livshits, B. (2019). The anatomy of a cryptocurrency {Pump-and-Dump} scheme. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 1609–1625.
- [Yang et al., 2016] Yang, W., Jin, L., Ni, H., and Lyons, T. (2016). Rotation-free online handwritten character recognition using dyadic path signature features, hanging normalization, and deep neural network. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 4083–4088. IEEE.
- [Yang et al., 2022] Yang, W., Lyons, T., Ni, H., Schmid, C., and Jin, L. (2022). Developing the path signature methodology and its application to landmark-based human action recognition. In *Stochastic Analysis, Filtering, and Stochastic Optimization: A Commemorative Volume to Honor Mark HA Davis’s Contributions*, pages 431–464. Springer.
- [You et al., 2017] You, J., Li, X., Low, M., Lobell, D., and Ermon, S. (2017). Deep gaussian process for crop yield prediction based on remote sensing data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- [Yun et al., 2012] Yun, K., Honorio, J., Chattopadhyay, D., Berg, T. L., and Samaras, D. (2012). Two-person interaction detection using body-pose features and multiple instance learning. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 28–35. IEEE.
- [Zabczyk, 2020] Zabczyk, J. (2020). *Mathematical Control Theory*. Springer.
- [Zaheer et al., 2017] Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. (2017). Deep sets. *Advances in neural information processing systems*, 30.
- [Zhu et al., 2016] Zhu, W., Lan, C., Xing, J., Zeng, W., Li, Y., Shen, L., and Xie, X. (2016). Co-occurrence feature learning for skeleton based action recognition using regularized deep lstm networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.

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>