



Convolutional signature for sequential data

Ming Min¹ · Tomoyuki Ichiba¹

Received: 3 August 2021 / Accepted: 12 March 2022 / Published online: 4 April 2022
© The Author(s) 2022

Abstract

Signature is an infinite graded sequence of statistics known to characterize geometric rough paths. While the use of the signature in machine learning is successful in low-dimensional cases, it suffers from the curse of dimensionality in high-dimensional cases, as the number of features in the truncated signature transform grows exponentially fast. With the idea of Convolutional Neural Network, we propose a novel neural network to address this problem. Our model reduces the number of features efficiently in a data-dependent way. Some empirical experiments including high-dimensional financial time series classification and natural language processing are provided to support our convolutional signature model.

Keywords Signature · Rough paths · Convolutional neural networks · Sequential data

Mathematics Subject Classification 60L10 · 62R07 · 68T50

JEL Classifications C630

1 Introduction

Multi-dimensional sequential data analysis is an important research area in Machine Learning, Financial Mathematics and many other areas. There are several methods of analyzing sequential data recently developed in deep learning,

Research in part was supported by the National Science Foundation grant NSF DMS-2008427.

✉ Ming Min
m_min@pstat.ucsb.edu

Tomoyuki Ichiba
ichiba@pstat.ucsb.edu

¹ Department of Statistics and Applied Probability, South Hall, University of California, Santa Barbara, CA 93106, USA

e.g., Recurrent Neural Network (RNN) (Cho et al. 2014), GRU (Cho et al. 2014), LSTM Hochreiter and Schmidhuber (1997) and Transformer (Vaswani et al. 2017). They have been successfully applied into a variety of important tasks in Data Science, such as natural language processing, financial time series and medical data analyses. Another mainstream approach to the sequential data analyses is Bayesian learning, mostly involved with Gaussian Process (GP) (Williams and Rasmussen 2006), where by pre-determined a priori distribution, it has advantage in quantifying uncertainty up to some extent. For example, Raissi et al. (2018) use GP to solve nonlinear partial differential equations with noisy boundary observations. More recently, a novel mathematical object, called *signature*, has been proposed and received more attention, to summarize information of sequential data, see (Boedihardjo et al. 2016; Chevyrev and Lyons 2016; Levin et al. 2013; Lyons and Qian 2002; Lyons et al. 2007). In this paper, we shall discuss the signature in the multi-dimensional sequential data analysis.

Signature is a graded feature set of a stream, or sequential data set, which is derived from the Rough Path Theory. Signature has been introduced as a feature map into the field of Machine Learning with successful applications to the sequential data. With truncations up to a given desirable accuracy, this special feature set has universality for approximations and can characterize pathwise data efficiently. It is known that the high-frequency sequential data set is transformed into several features efficiently by the truncated signature in the case of relatively low dimensional paths. For instance, Lyons and Oberhauser (2014) use the signatures to characterize high-frequency trading paths of financial data. Kidger et al. (2019) use the signature transform as a layer of neural network, and propose the deep signature transform model. Moreover, the use of the signature is model-free and data-dependent, see (Lyons et al. 2019, 2020).

However, the application of the signatures suffers from the curse of dimensionality, because the number of features in the truncated signature grows exponentially fast as the dimension increases. Consequently, in the case of high-dimensional sequential data, this feature map requires computational costs in real data analyses. A kernel-based learning algorithm has been introduced to address this problem in Kiraly and Oberhauser (2019); Toth and Oberhauser (2019). In this paper, we propose a new algorithm to solve this problem by combining Convolutional Neural Network and the signature transform. We evaluate the reduction of the number of features and show by numerical experiments that this algorithm can gain efficiency. Thus, this algorithm may contribute many applications of the signatures in the cases of the high-dimensional sequential data.

The rest of this paper is organized as follows. In Sect. 2, we review the signature of rough paths, geometric rough paths and nice properties, discuss Signature Classifier in the classification problems, as a typical application of the signature, and evaluate its classification error in Theorem 1. In Sect. 3, we introduce the main algorithm of this paper, a Convolutional Signature model, evaluate how this model reduces the number of features, show that this model preserves all information of path data and discuss its universality in Theorem 2. In Sect. 4, a broad

range of experiments are performed to support our model, including high-dimensional, financial time series classification, functional estimation and textual sentiment detection. We conclude with further ongoing research in Sect. 5.

2 Signature and geometric rough paths

2.1 Signatures

Let us introduce some notations for the sequential data sets, in order to explain the signature method, following (Lyons et al. 2007). Given a Banach space E with a norm $\| \cdot \|$, we define the tensor algebra

$$T((E)) := \{(a_i)_{i \geq 0} : a_i \in E^{\otimes i} \text{ for every } i\} \tag{2.1}$$

associated with the sum $+$ and with the tensor product \otimes defined by

$$(a_i)_{i \geq 0} + (b_i)_{i \geq 0} := (a_i + b_i)_{i \geq 0}, \quad (a_i)_{i \geq 0} \otimes (b_i)_{i \geq 0} := (c_i)_{i \geq 0},$$

where the j th element $c_j := \sum_{k=0}^j a_k \otimes b_{j-k}$ is the convolution of the first j elements of $(a_i)_{i \geq 0}$ and $(b_i)_{i \geq 0}$ in $T((E))$. Similarly, let us define its subset

$$T(E) := \{(a_i)_{i \geq 0} : a_i \in E^{\otimes i} \text{ and } \exists N \in \mathbb{N} \text{ such that } a_i = 0 \forall i \geq N\} \tag{2.2}$$

of $T((E))$ for those with finite number of non-zero elements. Note that $T(E) \subset T((E))$. Also, we shall consider the truncated tensor algebra of order $m \in \mathbb{N}$, i.e.,

$$T^m(E) := \{(a_i)_{i=0}^m : a_i \in E^{\otimes i} \text{ for } \forall i \leq m\}, \tag{2.3}$$

which is a subalgebra of $T((E))$. Then as we shall see, the signatures and the m -th order truncated signatures lie in these spaces $T((E))$ and $T^m(E)$, respectively.

Now with $E := \mathbb{R}^d$ and the usual Euclidean norm $\| \cdot \|$, we shall define the space $\mathcal{V}^p([0, T], E)$ of the d -dimensional continuous paths of finite p -th variation over the time interval $[0, T]$ and the signatures of the paths in $\mathcal{V}^p([0, T], E)$.

Definition 1 (*The space of finite p -variation paths*) Fix $p \geq 1$ and the interval $[0, T]$. The p -variation of a d -dimensional path $X : [0, T] \rightarrow E := \mathbb{R}^d$ is defined by

$$\|X\|_p := \left(\sup_{D_n \subset [0, T]} \sum_{i=0}^{n-1} \|X_{t_{i+1}} - X_{t_i}\|^p \right)^{1/p}.$$

Here, the supremum is taken over all the possible partitions of the form $D_n := \{t_i\}_{1 \leq i \leq n}$ of $[0, T]$ with $0 = t_0 < t_1 < \dots < t_n \leq T$, $n \geq 1$. X is said to be of finite p -variation, if $\|X\|_p < \infty$. We denote the set of continuous paths $X : [0, T] \rightarrow E$ of finite p -variation by $\mathcal{V}^p([0, T], E)$.

Table 1 The corresponding signatures for the smooth path $(t, (t - 2)^3)$ and the piece-wise linear path of the augmented AAPL adjusted price in Fig. 1, respectively

X	$(t, (t - 2)^3)$	AAPL
X^1	(4, 16)	(1, 65.52)
X^2	$\begin{pmatrix} 8 & 32 \\ 32 & 128 \end{pmatrix}$	$\begin{pmatrix} 0.5 & 31.17 \\ 34.00 & 2123.3 \end{pmatrix}$

We use the supremum norm $\|\cdot\|_\infty$ for continuous functions on $[0, T]$, i.e., $\|f\|_\infty := \sup_{x \in [0, T]} |f(x)|$. It can be shown that if we equip the space $\mathcal{V}^p([0, T], E)$ with the norm $\|X\|_{\mathcal{V}^p([0, T], E)} := \|X\|_p + \|X\|_\infty$, then $\mathcal{V}^p([0, T], E)$ is a Banach space. Now the signature and truncated signature are defined as follows.

Definition 2 (Signatures) The signature $S(X)$ of a path $X \in \mathcal{V}^p([0, T], E)$, $p \geq 1$ is defined by $S(X) := (1, X^1, X^2, \dots) \in T((E))$, where the k -th element

$$X^k := \int \dots \int_{0 < t_1 < \dots < t_n < T} dX_{t_1} \otimes \dots \otimes dX_{t_n} \in E^{\otimes k} \tag{2.4}$$

is the k -fold, iterated integral for $k \geq 1$, if the iterated integrals are well defined.

The truncated signature is naturally defined as $S^m(X) := (1, X^1, X^2, \dots, X^m) \in T^m(E)$ for every $m \geq 1$ including the 0-th term $S^0(X) = 1$.

Remark 1 The integrals in (2.4) depend on the nature of the paths. Here are some typical examples:

1. If X is of 1-variation path, then the integrals (2.4) of the signature can be understood as the Stieltjes integral;
2. If X is of p -variation path with $1 < p < 2$, then it can be defined in the sense of Young (e.g., see Lyons and Qian 2002).
3. If X is a Brownian motion, then we can use the Itô integral or the Stratonovitch integral. As we will explain later, when extending from a Brownian motion path or a semimartingale to a geometric rough path, we choose the Stratonovitch integral rather than the Itô integral.

Example 1 (Smooth paths and piece-wise linear paths) For $p \geq 1$ the path space $\mathcal{V}^p([0, T], E)$ contains the smooth functions and the piece-wise linear functions. We give the following two examples of paths in $\mathcal{V}^p([0, T], E)$, as shown in Fig. 1. In its left panel, we plot the smooth path $X_t = (t, (t - 2)^3), t \in [0, 4]$. In its right panel, we represent the discrete data: daily AAPL adjusted close stock price from Nov 28, 2016 to Nov 24, 2017 by interpolating the path linearly between each successive two days. The first 2 degree signatures X^1 and X^2 of these two paths in (2.4) are calculated and given in Table 1.

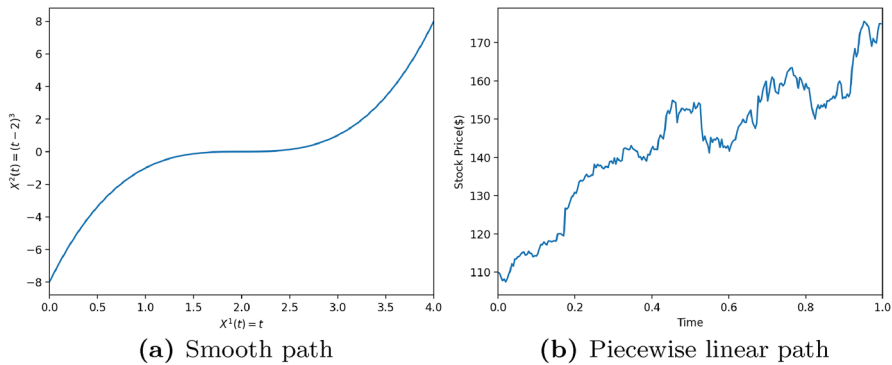


Fig. 1 Examples of $\mathcal{V}^p([0, T], E)$, $p \geq 1$: (a) Plot of a smooth path $X_t = (t, (t - 2)^3), t \in [0, 4]$. (b) Plot of linear interpolation of daily AAPL adjusted close stock price from Nov 28, 2016 to Nov 24, 2017

2.2 Geometric rough paths and linear functionals

Here we introduce rough paths and geometric rough paths briefly. More details can be found in Lyons and Qian (2002) and Lyons et al. (2007). Instead of $T((E))$ in (2.1), the p -rough paths and the geometric p -rough paths are objects in $T^{[p]}(E)$ in (2.3) for some real number $p (\geq 1)$. A fundamental result from rough paths theory and signatures (Lyons et al. 2007) is that there exists a continuous unique lift from $T^{[p]}(E)$ to $T((E))$. This lift is made in an iterated integral, and consequently, it gives us the signature of rough paths.

We denote the space of the p -rough paths by Ω_p . The space $G\Omega_p$ of the geometric p -rough paths is defined by the p -variational closure (cf. (Lyons et al. 2007) Chapter 3.2) of $S^{[p]}(\Omega_1)$. For a path $X : [0, T] \rightarrow \mathbb{R}^d$ with the bounded p -variation, the truncated signature belongs to the space of the p -rough paths, i.e., $S^{[p]}(X) \in \Omega_p$. If X is of bounded 1-variation, then the truncated signature belongs to the space of the geometric p -rough paths, i.e., $S^{[p]}(X) \in G\Omega_p$ for any $p (\geq 1)$.

It is manifested that the signature enjoys many nice properties. For example, signature characterizes paths up to tree-like equivalence (Boedihardjo et al. 2016) that are parametrization invariant. Here is a precise statement.

Proposition 1 (Parametrization Invariance, Lemma 2.12 of Levin et al. (2013)) *Let $X : [0, T] \rightarrow \mathbb{R}^d$ be a path with bounded variation and $\psi : [0, T] \rightarrow [0, T]$ a re-parametrization of the time parameter. If we define \tilde{X} by $\tilde{X}_t := X_{\psi(t)}$, then each term in $S(\tilde{X})$ is equal to the corresponding term in $S(X)$, i.e. $S(\tilde{X}) = S(X)$.*

Moreover, if there exists a monotone increasing dimension in the path with bounded variation or geometric rough path, we can get rid of tree-like equivalence (Boedihardjo et al. 2016; Gyurkó et al. 2013; Levin et al. 2013). Also, it is easy to specify one path among the parametrization invariance by adding timestamps. In other words, provided that an extra time dimension included, signature characterize geometric rough path uniquely. Another useful fact from rough path theory

(Chevyrev and Lyons 2016; Lyons and Qian 2002) is that signature terms enjoy a factorial decay as the depth increases, which makes truncating signature reasonable. The following remark shows an example of the factorial decay for bounded 1-variation paths.

Remark 2 (Factorial Decay, Proposition 2.2 of Lyons et al. (2007)) Let $X : [0, T] \rightarrow \mathbb{R}^d$ be a continuous path with bounded 1-variation, then for every $k \geq 1$

$$\left\| \int_{0 \leq t_1 < \dots < t_k \leq T} \dots \int dX_{t_1} \otimes \dots \otimes dX_{t_k} \right\| \leq \frac{\|X\|_1^k}{k!}, \tag{2.5}$$

where $\|\cdot\|$ is the tensor norm.

All these properties motivate us to use the signature as a feature map in Data Science. We shall then define the linear forms on the signatures.

For simplicity, let us fix $E = \mathbb{R}^d$, and let $\{e_i\}_{i=1}^d$ ($\{e_i^*\}_{i=1}^d$, respectively) be a basis of \mathbb{R}^d (a basis of the dual space $(\mathbb{R}^d)^*$ of \mathbb{R}^d , respectively). For every $n \in \mathbb{N}$ and indexes $(i_1, \dots, i_n) \in \{1, \dots, d\}^n$, $(e_{i_1}^* \otimes \dots \otimes e_{i_n}^*)$ can be naturally extended to $(E^*)^{\otimes n}$ with the basis $(e_I^* = e_{i_1}^* \otimes \dots \otimes e_{i_n}^*)$, and we call $I = i_1 \dots i_n$ a word of length n . The linear actions of $(E^*)^{\otimes n}$ on $E^{\otimes n}$ extends naturally a linear mapping $(E^*)^{\otimes n} \rightarrow T((E))^*$ by

$$e_I^*(\mathbf{a}) := e_I^*(a_n), \tag{2.6}$$

for every word I and every element $\mathbf{a} = (a_0, a_1, \dots, a_n, \dots) \in T((E))$.

Let A^* be the collection of all words of length n for all $n \in \mathbb{N}$. Then $\{e_I^*\}_{I \in A^*}$ forms a basis of $T(E^*) = T((\mathbb{R}^d)^*)$. Let $I, J \in A^*$ be two words of lengths m and n with $I = i_1 \dots i_m$ and $J = j_1 \dots j_n$, respectively. We say a permutation σ in the symmetric group \mathfrak{S}_{m+n} of $\{1, \dots, m+n\}$ is a shuffle of $\{1, \dots, m\}$ and $\{m+1, \dots, m+n\}$, if $\sigma(1) < \dots < \sigma(m)$ and $\sigma(m+1) < \dots < \sigma(m+n)$. We denote the collection of all shuffles of $\{1, \dots, m\}$ and $\{1, \dots, n\}$ by $Shuffles(m, n)$.

Definition 3 (Shuffle Product) For every pair $I = i_1 \dots i_m, J = j_1 \dots j_n$ of words of length m and n , the shuffle product $e_I^* \sqcup e_J^*$ of e_I^* and e_J^* is given by

$$e_I^* \sqcup e_J^* := \sum_{\sigma \in Shuffles(m, n)} e_{(k_{\sigma^{-1}(1)} \dots k_{\sigma^{-1}(m+n)})}^*, \tag{2.7} \quad \text{where}$$

$$k_1 \dots k_{m+n} = i_1 \dots i_m j_1 \dots j_n.$$

Denote $T((\mathbb{R}^d)^*)$ as the space of linear forms on $T((\mathbb{R}^d))$ induced by $T((\mathbb{R}^d)^*)$. The shuffle product between $f, g \in T((\mathbb{R}^d)^*)$ denoted by $f \sqcup g$ can be defined via natural extension of (2.7), by the bi-linearity of \sqcup . It can be shown that $T((\mathbb{R}^d)^*)$ is an algebra equipped with shuffle product and element-wise addition restricted to the geometric rough path space $S(\mathcal{V}^p([0, T], \mathbb{R}^d))$, see Theorem 2.15 of Lyons et al. (2007). The following proposition motivates us to use the signature as a feature map.

Proposition 2 (Universal Approximation) *Fix $p \geq 1$, a continuous function $f : \mathcal{V}^p([0, T], \mathbb{R}^d) \rightarrow \mathbb{R}$ of finite p -variation, and a compact subset K of $\mathcal{V}^p([0, T], \mathbb{R}^d)$. If $S(x)$ is a p -geometric rough path for each $x \in K$, then for every $\epsilon > 0$, there exists a linear form $l^\epsilon \in T((\mathbb{R}^d))^*$, such that*

$$\sup_{x \in K} |f(x) - \langle l^\epsilon, S(x) \rangle| < \epsilon. \tag{2.7}$$

Proof The proof follows directly from the uniqueness of signature transform for geometric rough paths and the Stone–Weierstrass theorem. See Lyons et al. (2020) and Theorem 4.2 in Arribas (2018) for more details. \square

Remark 3 (A curse of dimensionality) By Definition 2, the truncated signature $S^m(X)$ has a total of $\mathbf{d}_m := \sum_{k=0}^m d^k = (d^{m+1} - 1)/(d - 1)$ many terms for $m \geq 0$. The signature transform is an efficient feature reduction technique, when we have the d dimensional path sampled with high frequency in time. However, when the dimension d is large, the number of signature terms to be computed increases exponentially fast and makes the signature not easily applicable in practice.

To our best knowledge at this time, only Kiraly and Oberhauser (2019) and Salvi et al. (2021) introduce new algorithms of calculating the kernel of the signatures and Toth and Oberhauser (2019) discuss the application of the kernel methods to fix this high-dimensional problem. We introduce Convolutional Neural Network (CNN) to solve this problem in Sect. 3.

2.3 Classification via signature

Before we discuss the convolutional neural network in Sect. 3, we consider the application of the signatures to classification problems. In classification problems, we estimate the probability of an object belonging to each class. This estimation problem for the sequential data classification can be solved via the signature.

On a probability space, $(\Omega, \mathcal{F}, \mathbb{P})$ consider k classes, class 1, class 2, ..., class k , and n paired independent data $(x^i, y^i)_{1 \leq i \leq n}$, where each $x^i : [0, T] \rightarrow \mathbb{R}^d$ is the path datum and the corresponding label $y^i \in \{1, \dots, k\}$ is the class which x^i belongs to. We assume that the labels y^1, \dots, y^n are sampled from a common distribution and the conditional probability $\mathbb{P}(x^i \in \cdot | y^i)$ of x^i , given the class y^i , is a common probability distribution for $i = 1, 2, \dots, n$. Since we often observe the path dataset at discrete time stamps and we use piece-wise linear interpolations to connect among them, it is reasonable to assume that each path x in the dataset is of bounded 1-variation. Hence, its signature $S(x^i)$ is a geometric 1 rough path in Sect. 2.2.

Definition 4 (Classification problem) Our sequential classification problem is stated as follows: given training data $(x^i, y^i)_{1 \leq i \leq n}$, derive a classifier g for predicting the labels for unseen data (x, y) . Let $p_j(x) := \mathbb{P}(y = j|x)$ for $j = 1, \dots, k$. Our goal is to estimate these conditional probability $p_j(x)$ by $\hat{p}_j(x)$ for the path x of bounded

1-variation and classify x in the class $\arg \max_j \hat{p}_j(x)$ for $j = 1, \dots, k$ as accurate as possible.

Since the signature $S(x)$ of x determines the path x uniquely, it is reasonable to consider the signature $S(x)$ and a nonlinear continuous function $g : T((\mathbb{R}^d)) \rightarrow [0, 1]^k$, such that

$$g(S(x)) = (\hat{p}_1(x), \dots, \hat{p}_k(x))^T, \quad (2.8)$$

where \hat{p}_j 's are estimator of p_j 's, subject to $\sum_{j=1}^k \hat{p}_j(x) = 1$. Here, T represents the transpose of the vector.

For practical use, we use the truncate signature transforms, thanks to the factorial decay property (Remark 2) of the signature. With the truncation depth m , we obtain the estimate

$$g(S^m(x)) = (\hat{p}_1(x), \dots, \hat{p}_k(x))^T, \quad (2.9)$$

where $g : T^m(\mathbb{R}^d) \rightarrow [0, 1]^k$ is a nonlinear continuous function, and then the predicted label is given by

$$\hat{y} = \arg \max_j \hat{p}_j(x). \quad (2.10)$$

Definition 5 (Signature Classifier) We call $h : T((\mathbb{R}^d)) \rightarrow [0, 1]$ of the form (2.8) a signature classifier, where $T((\mathbb{R}^d))$ is the tensor algebra and h is a nonlinear continuous function. Naturally, a truncated signature classifier of degree $m \in \mathbb{N}$ is $h : T^m(\mathbb{R}^d) \rightarrow [0, 1]$ of the form (2.9).

In the simple case with only 2 classes, *class 0* and *class 1*, we consider the following concentration inequalities for classification via signature. We first restate the classification problem for the two classes. Suppose we have the pairwise, independent, identically distributed samples $(X^1, Y^1), \dots, (X^n, Y^n)$ where $Y^i \in \{0, 1\}$ and $X^i \in \mathcal{V}^1([0, T], \mathbb{R}^d)$. Let $h : \mathcal{V}^1([0, T], \mathbb{R}^d) \rightarrow \{0, 1\}$ be a classifier. The training error $\hat{R}_n(h)$ and the true error $R(h)$ are defined by

$$\hat{R}_n(h) = \frac{1}{n} \sum_{i=1}^n I(Y^i \neq h(X^i)), \quad \text{and} \quad R(h) = \mathbb{P}(Y \neq h(X)). \quad (2.11)$$

Here, $I(\cdot)$ is the indicator function. Correspondingly, $R(h) = \mathbb{P}(Y \neq I(h(X) > 0.5))$ and $\hat{R}_n(h) = \frac{1}{n} \sum_{i=1}^n I(Y^i \neq I(h(X^i) > 0.5))$. We shall see that $\hat{R}_n(\hat{h}) := \inf_{h \in \mathcal{H}} \hat{R}_n(h)$ is close to $R(h_*) := \inf_{h \in \mathcal{H}} R(h)$, where \mathcal{H} is the collection of the signature classifiers and we assume that $h_* \in \mathcal{H}$. Denote the set

$$\mathcal{E} := \{ \sup_{h \in \mathcal{H}} |\hat{R}_n(h) - R(h)| \leq \epsilon \}$$

to be the event that the training error $\hat{R}_n(h)$ is close to the true error $R(h)$ for all classifiers $h \in \mathcal{H}$ in the range of ϵ , given a fixed $\epsilon > 0$.

From now on, we assume \mathcal{H} is a compact set of truncated signature classifiers of degree m equipped with metric ρ . The following definition comes from van Handel (2016).

Definition 6 (δ -net and covering number) A set H is called a δ -net for (\mathcal{H}, ρ) if for every $h \in \mathcal{H}$, there exists $\pi(h) \in H$ such that $\rho(h, \pi(h)) < \delta$. The smallest cardinality of a δ -net for (\mathcal{H}, ρ) is called the covering number

$$N(\mathcal{H}, \rho, \delta) := \inf\{|H| : H \text{ is a } \delta\text{-net for } (\mathcal{H}, \rho)\}. \tag{2.12}$$

In our case, we may take the uniform norm ρ , for example. Indeed, by the Ascoli-Arzelà theorem, we only need \mathcal{H} to be equicontinuous to make it compact, and hence $N_\delta := N(\mathcal{H}, \rho, \delta)$ is always finite for any $\delta > 0$. Let H_δ be a δ -net of \mathcal{H} with cardinality N_δ .

Theorem 1 For every $\epsilon > 0$, $\epsilon_0 > 0$, there exist $\delta > 0$ and a corresponding finite covering number N_δ , such that

$$\mathbb{P}(\sup_{h \in \mathcal{H}} |\hat{R}_n(h) - R(h)| > \epsilon) \leq 2N_\delta e^{-2n\epsilon} + \epsilon_0. \tag{2.13}$$

Proof Take a δ -net H_δ of \mathcal{H} with cardinality N_δ . By the Markov inequality and the definition of the covering number, we have

$$\begin{aligned} \mathbb{P}(\sup_{h \in H_\delta} (\hat{R}_n(h) - R(h)) > \epsilon) &\leq e^{-t\epsilon} \mathbb{E}[\sup_{h \in H_\delta} e^{t(\hat{R}_n(h) - R(h))}] \\ &\leq N_\delta e^{-t\epsilon} \sup_{h \in H_\delta} \mathbb{E}[e^{t(\hat{R}_n(h) - R(h))}]. \end{aligned}$$

Since $\hat{R}_n(h)$ is the sum (2.11) of independent random variables, by Hoeffding’s inequality (Hoeffding 1963), we have $e^{-t\epsilon} \mathbb{E}[e^{t(\hat{R}_n(h) - R(h))}] \leq e^{-2n\epsilon}$ for $h \in H_\delta$, $t \geq 0$ and $n \geq 1$. Hence, for every $n \geq 1$ and δ -net H_δ of \mathcal{H} , we have

$$\mathbb{P}(\sup_{h \in H_\delta} (\hat{R}_n(h) - R(h)) > \epsilon) \leq N_\delta e^{-t\epsilon} \sup_{h \in H_\delta} \mathbb{E}[e^{t(\hat{R}_n(h) - R(h))}] \leq N_\delta e^{-2n\epsilon}.$$

By a similar argument, we also have $\mathbb{P}(\sup_{h \in H_\delta} (R(h) - \hat{R}_n(h)) > \epsilon) \leq N_\delta e^{-2n\epsilon}$ for every $n \geq 1$ and δ -net H_δ of \mathcal{H} .

Combining the above two inequalities, we obtain that for every $n \geq 1$ and δ -net H_δ of \mathcal{H}

$$\mathbb{P}(\sup_{h \in H_\delta} |\hat{R}_n(h) - R(h)| > \epsilon) \leq 2N_\delta e^{-2n\epsilon}.$$

By approximating the supremum over \mathcal{H} by the supremum over the sets H_δ with cardinality N_δ , that is,

$$\mathbb{P}(\sup_{h \in \mathcal{H}} |\hat{R}_n(h) - R(h)| > \epsilon) = \lim_{\delta \rightarrow 0} \mathbb{P}(\sup_{h \in H_\delta} |\hat{R}_n(h) - R(h)| > \epsilon),$$

we conclude (2.13) that for any $\epsilon_0 > 0$, there exists a $\delta > 0$,

$$\begin{aligned} \mathbb{P}(\sup_{h \in \mathcal{H}} |\hat{R}_n(h) - R(h)| > \epsilon) &< \mathbb{P}(\sup_{h \in H_\delta} |\hat{R}_n(h) - R(h)| > \epsilon) + \epsilon_0 \\ &\leq 2N_\delta e^{-2n\epsilon} + \epsilon_0. \end{aligned}$$

□

By Theorem 1, the event \mathcal{E} holds with high probability provided that n is sufficiently large. On the set \mathcal{E} , we have by definitions

$$R(h_*) \leq R(\hat{h}) \leq \hat{R}_n(\hat{h}) + \epsilon \leq \hat{R}_n(h_*) + \epsilon \leq R(h_*) + 2\epsilon. \tag{2.14}$$

Thus, it follows that $|R(\hat{h}) - R(h_*)| \leq 2\epsilon$ on the set \mathcal{E} . Thus, on \mathcal{E} , the best empirical signature classifier \hat{h} is close to the best true signature classifier h_* as in (2.14). The connection between signature classifier and general classifier can be constructed by the uniqueness of the signature transform.

This covering number N_δ in Definition 6 plays an essential role here. The study of the covering number $N(\mathcal{H}, \rho, \delta)$ for the compact set \mathcal{H} of the truncated signature classifiers is still in progress. If we can quantify this number, then the number of training samples n needed for fixed error can be calculated from (2.13).

Example 2 (GARCH time series) We give an example of two classes of time series, $\{x^n\}_{n=1}^N$, generated by GARCH(2,2) model. The time series are given by

$$\begin{aligned} x_k^n &= \sigma_k \epsilon_k, \\ \sigma_k^2 &= w + \sum_{i=1}^2 \alpha_i x_{k-i}^n + \sum_{j=1}^2 \beta_j \sigma_{k-j}^2, \end{aligned}$$

where $w > 0$, $\alpha_i \geq 0$, $\beta_j \geq 0$ and ϵ_k 's are I.I.D. standard normal distributed. Denote $\alpha = (\alpha_1, \alpha_2)$ and $\beta = (\beta_1, \beta_2)$. 2 classes of GARCH time series are generated by setting parameters in Table 2.

For paths x^n generated by the first row parameters in Table 2, we label $y^n = 1$ (class 1), for the rest paths x^n generated by the second row parameters in Table 2, we label them by $y^n = 2$ (class 2). Thus, we generate paired data $\{(x^n, y^n)\}_{n=1}^N$.

Table 2 Parameters for GARCH(2,2) time series

Class	w	α	β
1	0.5	(0.4, 0.1)	(0.7, 0.5)
2	0.2	(0.8, 0.5)	(0.4, 0.1)

Remark 4 It is important to note that we cannot directly apply Proposition 2 here, because this $p(x)$ may not be continuous in x . Intuitively, it is better to add nonlinearity on classifier $h(\cdot)$. The experiment in Sect. 4.2 verifies this intuition.

In practice, the signature classifier (2.9) and its truncation (2.10) can be applied to find the classification model $g(\cdot)$ to estimate \hat{y} in other contexts. In Sect. 4, we shall apply the logistic regression to Example 2, and the result shows that the use of the truncated signature to classify this GARCH(2,2) time series is significantly efficient.

3 Convolutional Signature Model

The main goal of this section is to introduce the Convolutional Signature (CNN-Sig) model. As we have seen in Remark 3 in Sect. 2.2, the truncated signature suffers from the exponential growth of the number \mathbf{d}_m of terms, when the dimension d is large, and in this case both space and time complexity increase dramatically. We will use Convolutional Neural Network (CNN) to reduce this exponential growth to at most linear growth. CNN has been mostly used in analyzing visual imagery, where it takes advantage of the hierarchical patterns in image and assembles complex patterns by focusing on many small pieces of the picture. Convolutional layer convolves the input data with a small rectangular kernel, and the output data can be masked with an activation function. As there are some patterns between channels of a path, this motivates us to consider the signature with CNN to address the high dimensional problem.

Before introducing the CNN-Sig model, we shall explain that the signature transform can be viewed as a layer in the deep neural network model.

3.1 Signature as a layer

Signature transform can be viewed as a layer in deep neural networks and this is firstly proposed in Kidger et al. (2019). In the background of Python package **signatory** Kidger and Lyons (2020), signature transform takes input tensor of shape (b, n, d) , corresponding to a batch of size b of paths in \mathbb{R}^d with n observing points at times $\{t_j\}_{j=1}^n$, and returns a tensor of shape (b, \mathbf{d}_m) or a stream like tensor of shape (b, n, \mathbf{d}_m) , where \mathbf{d}_m is defined in Remark 3 in Sect. 2.2. Usually it omits the first term 1 of the signature transform. Since the signature is also differentiable

numerically with respect to each data point, the backpropagation calculation is available. In this way, the signature can be viewed as a layer in neural network.

3.2 Convolutional signature model

CNN, which has been proved to be a powerful tool in computer vision, is an efficient feature extraction technique. This idea has been used in Liao et al. (2019) as well as the “**Augment**” module (Kidger and Lyons 2020) (but only 1D CNNs are used). There are two cases of using 1D CNNs. The first case is to extract new sequential features of original paths and then paste them to the original path as extra dimensions. This method is not helpful in the high dimensional case and causes extra difficulty. The second case is that we use extracted sequential features directly from the 1D CNN. It works as a dimension reduction technique but the challenge is that it causes loss of information.

With the favor of the 2D CNN, we are able to reduce the number of signature features and capture all information in the original path at the same time. Since the convolution here is different from the convolution concept in mathematics, we define it and present Example 3 to show the computational details for those who are not so familiar with CNN.

Definition 7 (2D Convolution) Let $*$ be an operation of element-wise matrix multiplication and summation between two matrices of the same shape, that is, $A := (a_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$ and $B := (b_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$ of the same size: $A * B = \sum_{i=1}^m \sum_{j=1}^n a_{ij} b_{ij}$. Suppose the input tensor is $M := (M_{ij})_{1 \leq i \leq l, 1 \leq j \leq l}$, a kernel window $K := (k_{ij})_{1 \leq i \leq m, 1 \leq j \leq n}$ and a stride window (s, t) . The output $O := (o_{p,q})$ of 2D convolution is given by

$$o_{p,q} := (M_{ij})_{1+(p-1)s \leq i \leq m+(p-1)s, 1+(q-1)t \leq j \leq n+(q-1)t} * K. \tag{3.1}$$

The shape of the output O depends on how we treat the boundary specifically and does not play a crucial role here.

Example 3 (2D Convolution) Let us consider a tensor $M := (M_{ij})_{1 \leq i,j \leq 5}$ and a kernel window $K := (k_{ij})_{1 \leq i,j \leq 3}$,

$$M := \begin{pmatrix} 2 & 1 & 0 & 2 & 0 \\ 0 & 1 & 2 & 2 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 2 & 0 & 0 & 2 & 2 \\ 0 & 2 & 0 & 1 & 1 \end{pmatrix}, \quad \text{and} \quad K := \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ -1 & -1 & -1 \end{pmatrix}, \text{ respectively,}$$

and a stride window $(1, 1)$. The output will be a 3×3 tensor, denoted by $O = (o_{ij})_{1 \leq i,j \leq 3}$, where each element o_{ij} of O is given by the element-wise multiplication and summation of




Fig. 2 Convolutional neural network and signature transform connected by Φ

$$\tilde{M}^{i,j} := (M_{k,\ell})_{i \leq k \leq i+2, j \leq \ell \leq j+2}$$

and K , i.e., $o_{i,j} = \tilde{M}^{i,j} * K$ for $1 \leq i, j \leq 3$. For example,

$$o_{11} = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 1 & 2 \\ 0 & 0 & 0 \end{pmatrix} * \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ -1 & -1 & -1 \end{pmatrix} = 2 \cdot 0 + 1 \cdot 1 + 0 \cdot 0 + \dots + 0 \cdot (-1) = -1,$$

$$o_{12} = \begin{pmatrix} 1 & 0 & 2 \\ 1 & 2 & 2 \\ 0 & 0 & 1 \end{pmatrix} * \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ -1 & -1 & -1 \end{pmatrix} = 1 \cdot 0 + 0 \cdot 1 + 2 \cdot 0 + \dots + 1 \cdot (-1) = -2,$$

$$o_{13} = \begin{pmatrix} 0 & 2 & 0 \\ 2 & 2 & 1 \\ 0 & 1 & 1 \end{pmatrix} * \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ -1 & -1 & -1 \end{pmatrix} = 1, \quad o_{21} = \begin{pmatrix} 0 & 1 & 2 \\ 0 & 0 & 0 \\ 2 & 0 & 0 \end{pmatrix} * \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & -1 \\ -1 & -1 & -1 \end{pmatrix} = -1,$$

and so on. Therefore, the output O is given by

$$O = \begin{pmatrix} -1 & -2 & 1 \\ -1 & -1 & -1 \\ 0 & -5 & -3 \end{pmatrix}.$$

The Convolutional Signature model uses the 2D CNN before the signature transform, and the structure of the convolutional signature model can be described in Fig. 2. The convolution is implemented in channels. Since the signature is efficient in the time direction, we do not have to convolute the time direction.

3.3 Number of features

Suppose $c (\leq d)$ is an integer such that d is divisible by c and let us fix the ratio $\gamma = d/c \in \mathbb{N}$. For the sake of simplicity of explanations, we set the number of features with kernel window of size $(1 \times c)$ and stride $(1 \times c)$. We illustrate our idea in the following example.

Example 4 Let us consider a tensor $M := (M_{i,j})_{1 \leq i \leq 5, 1 \leq j \leq 4}$ and 2 kernel windows $K_1 := (k_i^1)_{1 \leq i \leq 2}, K_2 := (k_i^2)_{1 \leq i \leq 2}$,

$$M := \begin{pmatrix} 2 & 1 & 0 & 2 \\ 0 & 1 & 2 & 2 \\ 0 & 0 & 0 & 1 \\ 2 & 0 & 0 & 2 \\ 0 & 2 & 0 & 1 \end{pmatrix}, \quad K_1 := (-1 \ 1) \quad \text{and} \quad K_2 := (1 \ 2).$$

Using a stride window $(1, 2)$, we calculate the output $O = \{O_1, O_2\}$ with $O_l = (o_{i,j}^l)_{1 \leq i \leq 5, 1 \leq j \leq 2}, l = 1, 2$. The computation is done in the same way as in Example 3: $o_{1,1}^1 = (2, 1) * (-1, 1) = -2 + 1 = -1, o_{2,2}^1 = (2, 2) * (-1, 1) = -2 + 2 = 0, o_{1,1}^2 = (2, 1) * (1, 2) = 2 + 2 = 4, o_{2,2}^2 = (2, 2) * (1, 2) = 2 + 4 = 6$. Therefore, the output O is given by

$$O_1 = \begin{pmatrix} -1 & 2 \\ 1 & 0 \\ 0 & 1 \\ -2 & 2 \\ 2 & 1 \end{pmatrix}, \quad O_2 = \begin{pmatrix} 4 & 4 \\ 2 & 6 \\ 0 & 2 \\ 2 & 4 \\ 4 & 2 \end{pmatrix}.$$

In this example, since K_1 and K_2 are linear-independent, we fully recover the input M given K_1, K_2 and output O .

Notice that since the first term in signature transform is always 1, we can omit that, to save the computational memory. As shown in Fig. 2, we start from one d -dimensional path with length L , using such a convolutional layer, and we resulted in c paths with each of d/c -dimensional. Then we augment each path with extra time dimension and apply signature transform to each path truncated at depth m , which gives us the number of features

$$N_f := c \cdot \frac{(d/c + 1)^{m+1} - d/c - 1}{d/c + 1 - 1} = \frac{(\gamma + 1)^{m+1} - \gamma - 1}{\gamma^2} \cdot d \tag{3.2}$$

many features by concatenating all c filters. These features can be used in any following neural network model. For example, a fully connected neural network in the simplest case, or a recurrent neural network (RNN) if we compute the sequence of the signature transform.

The number N_f of features grows linearly in d by increasing c linearly and fixing γ . Instead of optimizing this N_f by setting $\gamma = \arg \min N_f$ directly, we can

think γ as a hyperparameter to be tuned to avoid overfitting problem. It can be easily seen that by setting $\gamma = 1$, we reach a minimum of N_f when $m \geq 3$. However, lower γ will give us higher c , which increase the number of parameters in the CNN step. We consider the sum $N_f + (\frac{d}{\gamma})^2$ of number of features and the number of parameters in CNN. Moreover, we can add a multiplier α to the second term, and then define a regularized number on γ ,

$$N^\alpha(\gamma) := \frac{(\gamma + 1)^m - 1}{\gamma^2} \cdot (\gamma + 1) \cdot d + \alpha \cdot \frac{d^2}{\gamma^2}. \tag{3.3}$$

We can select a large real positive number α . This will help us avoid the overfitting problem, when we are concerned about that the CNN layer fits the original paths too well and it sacrifices the prediction power.

3.4 One-to-one Mapping

Under the setup in Sect. 3.3, we can generalize Example 4 and prove such a convolutional layer preserves all information of the original path. Suppose that $\{k^i\}_{i=1}^c$ are all c convolutional kernels with $k^i = (k_1^i, \dots, k_c^i)$ for $i = 1, \dots, c$. Denote the square matrix

$$\mathbf{K} := \begin{pmatrix} k_1^1 & \dots & k_c^1 \\ \vdots & \vdots & \vdots \\ k_1^c & \dots & k_c^c \end{pmatrix}.$$

Let the original path be $\mathbf{x} = (x_{t_1}, \dots, x_{t_n})^\top$, $x_{t_j} = (x_{t_j}^1, \dots, x_{t_j}^d)$ and the output path $\{\tilde{x}_i\}_{i=1}^c$, where $\tilde{x}_i = (\tilde{x}_{t_1,i}, \dots, \tilde{x}_{t_n,i})^\top$ with $\tilde{x}_{t_j,i} = (\tilde{x}_{t_j,i}^1, \dots, \tilde{x}_{t_j,i}^\gamma)$, $1 \leq i \leq c$. The CNN layer can be represented in equation as

$$\mathbf{K} \cdot (x_{t_j}^{lc+1}, \dots, x_{t_j}^{(l+1)c})^\top = (\tilde{x}_{t_j,1}^l, \dots, \tilde{x}_{t_j,c}^l)^\top, \quad 1 \leq l \leq \gamma, \quad 1 \leq j \leq n. \tag{3.4}$$

Lemma 1 *If \mathbf{K} is of full rank, then this CNN layer is a one-to-one map.*

Proof Since \mathbf{K} is square and of full rank, it is invertible.

$$(x_{t_j}^{lc+1}, \dots, x_{t_j}^{(l+1)c})^\top = \mathbf{K}^{-1} \cdot (\tilde{x}_{t_j,1}^l, \dots, \tilde{x}_{t_j,c}^l)^\top, \quad 1 \leq l \leq \gamma, \quad 1 \leq j \leq n.$$

It follows that the original path \mathbf{x} can be fully recovered by $\tilde{x} := \{\tilde{x}_i\}_{i=1}^c$. □

We denote the CNN layer transform as $\mathbf{K} : \mathcal{V}^1([0, T], \mathbb{R}^d) \rightarrow \mathcal{V}^1([0, T], \mathbb{R}^{d/c+1})^c$. Here, plus 1 in the dimension $(d/c) + 1$ comes from the time dimension we add to each convoluted path.

In accordance with practical case, we consider approximating functions with domain in a subspace of $\mathcal{V}^1([0, T], \mathbb{R}^d)$ that is observed at finite time stamps and connected by linear interpolation between consecutive points. More precisely, define

$$\begin{aligned} \mathcal{V}_D^1([0, T], \mathbb{R}^d) &:= \{x \in \mathcal{V}^1([0, T], \mathbb{R}^d) : \text{there exist } n \in \mathbb{N} \text{ and } 0 = t_0 < \dots < t_n = T \\ &\text{such that } x(t) = \frac{t_i - t}{t_i - t_{i-1}}x(t_{i-1}) + \frac{t - t_{i-1}}{t_i - t_{i-1}}x(t_i) \\ &\text{for } t_{i-1} \leq t \leq t_i, i = 1, \dots, n\}. \end{aligned} \quad (3.5)$$

Suppose $f : \mathcal{V}_D^1([0, T], \mathbb{R}^d) \rightarrow \mathbb{R}$ is the continuous function we need to estimate. Then we have the following theorem.

Theorem 2 (Approximation by the CNN-Sig model) *Let K be a compact set in $\mathcal{V}_D^1([0, T], \mathbb{R}^d)$. Suppose that f is Lipschitz in K . For any $\epsilon > 0$ there exist a CNN layer \mathbf{K} , an integer m , and a neural network model Φ such that*

$$\sup_{x \in K} |f(x) - \Phi \circ S^m \circ \mathbf{K}(x)| < \epsilon.$$

Proof For every $x \in \mathcal{V}_D^1([0, T], \mathbb{R}^d)$, we rewrite $f(x)$ as a function of $\tilde{x} = \{\tilde{x}_i\}_{i=1}^c$ in (3.4):

$$f(x) = f(\mathbf{K}^{-1}(\tilde{x})) = f \circ \mathbf{K}^{-1}(\tilde{x}) =: h(\tilde{x}). \quad (3.6)$$

It follows that $h = f \circ \mathbf{K}^{-1}$ is a continuous function. Since $S(\tilde{x}_i)$ is a geometric rough path and characterize the path \tilde{x}_i uniquely for each $1 \leq i \leq c$, there exists a continuous function $\hat{h} : (T(\mathbb{R}))^c \rightarrow \mathbb{R}$ such that

$$h(\tilde{x}) = \hat{h}(S(\tilde{x}_1), \dots, S(\tilde{x}_c)).$$

The existence follows from the compactness and that the signature map is continuous and one-to-one. Moreover, since f is Lipschitz, we have that h is Lipschitz and hence \hat{h} is also Lipschitz. The compactness of K implies that the image of $S \circ \mathbf{K}$ is also compact, hence $h(\tilde{x})$ can be approximate arbitrarily well by truncated signatures up to a uniform truncation depth m for all data in the set K . The existence of such m is induced by the proof of (Lemma 4.1, Min and Hu (2021)) and Lipschitz property. That is, there exists an integer m , such that

$$\sup_{x \in K} |\hat{h}(S(\tilde{x}_1), \dots, S(\tilde{x}_c)) - \hat{h}(S^m(\tilde{x}_1), \dots, S^m(\tilde{x}_c))| \leq \frac{\epsilon}{2}. \quad (3.7)$$

This \hat{h} is not necessarily linear, because there might be some dependence among $\{\tilde{x}_i\}_{i=1}^c$, but it can be approximated by a neural network model arbitrarily well. A wide range of Φ can be chosen. For example, a fully connected shallow neural network with one wide enough hidden layer and some activation function would work, see Funahashi (1989), Cybenko (1989); or a narrow but deep network, see Kidger and Lyons ((2020)). That is, there exists Φ such that

$$\sup_{x \in K} \left| \Phi(S^m(\tilde{x}_1), \dots, S^m(\tilde{x}_c)) - \hat{h}(S^m(\tilde{x}_1), \dots, S^m(\tilde{x}_c)) \right| \leq \frac{\epsilon}{2}. \quad (3.8)$$

By combining (3.6), (3.7), (3.8) together, we get the desired result. \square

In the CNN-Sig model, the CNN layer can be understood as data-dependent encoder which help us find the best way of encoding original path to several lower-dimensional paths. On the one hand, a large c will result in overfitting problem of CNN layer. On the other hand, small c will produce large number of features for Φ , and then Φ may have the overfitting problem. This tradeoff can be balanced by minimizing $N^\alpha(\gamma)$ in Eq. (3.3). Thus, although the choice of c does not affect the universality of the model, it could help with resolving the overfitting problem.

Remark 5 When we do experiments of the CNN-Sig model, this model works even better compared to plain signature transform of original path on testing data, it is because the CNN-Sig model reduces the number of features and thus overcomes the overfitting problem better than direct signature transform.

Moreover, the signature transform can be performed in a sequential way. Then we can choose a RNN model (GRU or LSTM) for Φ . Some other candidates for Φ can be Attention model like Transformer, 1d-CNN and so on, which might help us get better predictions. Thus, this CNN-Sig model is quite flexible and can be incorporated with many other well-developed deep learning models as Φ , which depends specifically on the task. In practice, we can use a different stride size to allow some overlap during convolution and reduce the number of filters. The one-to-one mapping property may be lost in this case if we choose a small number of filters, but it results in less overfitting. Another alternative is that we can also convolute over time dimension, provided that correlation over time is of importance to the sequential data.

4 Experiments

In this section, several results of the experiments are provided for the purpose of exhibiting the performance of the signature classifier and the CNN-Sig model. Sections 4.1 and 4.2 show that the signature classifier can be a nice candidate for the time series classification problem. In Sects. 4.3 and 4.4, we apply the CNN-Sig model to high-dimensional tasks, including the standard high-dimension datasets, approximation of maximum-call European payoff and sentimental analysis.

Table 3 Training (left) and testing (right) confusion matrices

True	0	1	True	0	1
Predicted			Predicted		
0	343	7	0	147	3
1	18	332	1	6	144

4.1 Classification of GARCH time series

The generalized autoregressive conditional heteroskedasticity (GARCH) process is usually used in econometrics to describe the time-varying volatility of financial time series (Bollerslev 1986; Engle 1982). GARCH provides a more real-world context than other models when predicting the financial time series, compare to other time series model like ARIMA. We apply logistic regression to Example 2, i.e. the goal is to estimate $g(S^m(x)) = (\hat{p}_0, \hat{p}_1)$ in (2.9), where

$$\log \frac{\hat{p}_1}{1 - \hat{p}_1} = \langle l, S^m(x) \rangle, \tag{4.1}$$

subject to $\hat{p}_0 + \hat{p}_1 = 1$, l is a linear functional on $T^m(\mathbb{R}^d)$ to be chosen such that the cross entropy

$$E(l) = - \sum_{i=1}^N (y^i \log \hat{p}_i + (1 - y^i) \log(1 - \hat{p}_i)) \tag{4.2}$$

is minimized, and we predict labels by $\hat{y}^i = \arg \max_i \hat{p}_i$. 500 samples are generated for each class and we use 70% of each class as training data and 30% of each as testing data. Using $m = 4$, we get training accuracy 96.4% and testing accuracy 97.0%. The confusion matrix is given below in Table 3.

4.2 Classification of directed chain discrete time series

In the study of mean-field interaction and financial systemic risk problems, (Detering et al. 2021) propose a countably many particle system of diffusion processes, coupled through an infinite, chain-like directed graph, and discuss a detection problem of mean-field interactions among diffusive particles. In Remark 4.5 of Detering et al. (2021), a discrete time analogue of the mean-reverting diffusions on the directed chain is also proposed.

We shall discuss a classification problem of such time series data partially observed from the directed chain graph. More specifically, we analyze an identically distributed time series data $\{X_n\}_{n \geq 1}$ and $\{\tilde{X}_n\}_{n \geq 1}$ parametrized by $a, u \in [0, 1]$ and defined recursively by

$$X_n = aX_{n-1} + (1 - a)(u\tilde{X}_{n-1} + (1 - u)\mathbb{E}[X_{n-1}]) + \varepsilon_n, \quad n \geq 1, \tag{4.3}$$

Table 4 Training accuracy and testing accuracy on Logistic regression

Training Acc	Testing Acc
0.7465	0.7375

Table 5 Training accuracy and testing accuracy on NN

Training Acc	Testing Acc
0.8930	0.8925

where we assume that $X_0 = \tilde{X}_0 = 0$ for simplicity, the distribution of $\{X_n, n \geq 0\}$ is identical with that of $\{\tilde{X}_n, n \geq 0\}$ and $\epsilon_n, n \geq 1$ are independent, identically distributed standard normal random variables, independent of $\{\tilde{X}_n\}_{n \geq 1}$. The parameter $u \in [0, 1]$ measures how much X_n depends on its neighborhood and $1 - u$ measures how much X_n depends on the common distribution. X and \tilde{X} have the same distribution with the moving average representation:

$$\begin{aligned}
 X_n &= \sum_{0 \leq l \leq k \leq n-1} \binom{k}{l} u^l (1-u)^{k-l} \epsilon_{n-k,l}, \\
 \tilde{X}_n &= \sum_{0 \leq l \leq k \leq n-1} \binom{k}{l} u^l (1-u)^{k-l} \epsilon_{n-k,l+1}, \quad n \geq 1,
 \end{aligned}
 \tag{4.4}$$

where $\{\epsilon_{n,k}, n, k \geq 0\}$ is an independent, identically distributed array of standard normal random variables.

Suppose that our only observation is $\{X_n\}_{n \geq 1}$, but both $\{\tilde{X}_n\}_{n \geq 1}$ and u are hidden to us. Our question is that given the access to $\{X_n\}_{n \geq 1}$ generated by different u , can we determine their classes?

In this part, we first set the default parameters and generate training and testing paths according to (4.4). First, we initial some parameters: $a = 0.5, u = 0.2$ or 0.8 for classification task, $N = 100$ is the time steps, $1/N$ is the variance of ϵ . To generate paths, we generate a $n \times (n + 1)$ matrix \mathcal{E} of the error terms ϵ , and then pick the column we need for each n . The summation takes time $O(N^2)$ and

Table 6 Training and testing mean absolute error(MAE) and R^2 for the direct signature transform plus linear regression (Sig+LR) and the CNN-Sig model with Φ as a fully connected neural network

d	Sig+LR				CNN-Sig			
	Training		Testing		Training		Testing	
	MAE	R^2	MAE	R^2	MAE	R^2	MAE	R^2
6 ($\gamma = 2$)	0.001	1.000	0.101	0.538	0.020	0.986	0.030	0.972
10 ($\gamma = 2$)	0.000	1.000	0.124	0.806	0.033	0.988	0.062	0.962
12 ($\gamma = 2$)	0.000	1.000	0.153	0.821	0.048	0.981	0.111	0.924
20 ($\gamma = 1$)	0.000	1.000	0.225	0.838	0.177	0.916	0.203	0.892

Fig. 3 QQ plot for training and testing result for lower dimensional regression with $d = 6, 10, 12, 20$ using the CNN-Sig model

we have to range n from 1 to N . The time complexity is the order of $O(N^3)$. We simulate 2000 training paths and 400 testing paths for this task.

Method 1: Logistic Regression In this method, we use 2000 training paths: 1000 for $u = 0.2$ and 1000 for $u = 0.8$. Calculating the signature transform of these paths, augmented with time dimension, up to degree 9, we build a Logistic Regression model on the signatures of training data and test this model, see Eq. (4.1).

The result is shown in Table 4. We observe that signature does capture useful features for u in these special time series.

Method 2: Deep Neural Network We build a Neural Network model to get a better result. We use 4 hidden layers with 256, 256, 128, 2 units respectively. For first 3 layers, we use "ReLU" as activation function, for last layer, we use "Softmax" activation function as the approximated probability values. After training for 20 epochs, the result is shown in Table 5.

This 4-layer neural network model produces better accuracy than logistic regression. The reason follows Remark 4. Logistic regression trains a linear classifier, but it cannot be used to estimate $p(\cdot)$ efficiently, because $p(\cdot)$ is not continuous in x . This DNN model adds nonlinearity to $h(\cdot)$ s and hence works better.

4.3 High-dimensional time series

Signature is an efficient tool as a feature map for high-frequency sequential data to reduce the number of features. However, the number of signature terms increases exponentially as dimension (or channels in the language of PyTorch) increasing. In Sect. 3, we proposed the CNN-Sig model to address this problem. We test our model by applying it in both regression and classification problem.


Experiments—Regression Problem for Maximum-Call Payoff

We investigate our model on a specific rainbow option, high-dimension European type maximum call option. In other words, we want to use our CNN-Sig model to estimate the payoff


$$\max_{1 \leq k \leq d} ((X_T^k - K)^+),$$

where T is terminal time, K is strike price, superscript k represents the k -th coordinate of this d -dimension path. If X_T^k is smaller than K for all $1 \leq k \leq d$, this payoff is zero. Otherwise the payoff would be the maximum of $(X_T^k - K)$ over those k satisfies $X_T^k \geq K$. Result of this experiment may motivate us to use CNN-Sig model in high-dimensional optimal stopping problem from financial mathematics.


Because of the limitation of exponential growth in the number of features, we use lower $d = 6, 10, 12, 20$ to compare the performance between plain signature




(a) train d=6




(b) test d=6




(c) train d=10




(d) test d=10




(e) train d=12



(f) test d=12



(g) train d=20



(h) test d=20

transform and CNN-Sig model. Then we apply this model to test its performance with higher dimension $d = 50$.

We generate 1000 training paths and 1000 testing paths for cases of $d = 6, 10, 12$, and generate 3000 training paths and 1000 testing paths for case $d = 20$. All stock price paths follows Black–Scholes model.

For all 4 cases, we consider $m = 4$ as the signature depth. For Φ in the CNN-Sig model, we use the same structure, 2 fully connected layers followed by ReLu activation function and then a fully connected layer. We did not apply any technique for avoiding overfitting problem in the CNN-Sig model to make this comparison fair. The result for comparison is shown in Table 6. We can see that for all these 4 cases, the CNN-Sig model beats direct signature transform. Since the CNN-Sig model reduces the number of features, it can help avoid overfitting problem compare to Sig+LR. We produce the QQ plots for training and testing results of the CNN-Sig model, see Fig. 3.

For $d = 50$, where the plain Sig+LR becomes not applicable, we use the same CNN-Sig structure as lower d cases for training. The training MAE is 0.206 with $R^2 = 0.982$ and testing MAE is 0.751 with $R^2 = 0.797$. The QQ plot of training and testing results is in Fig. 4. In this experiment, we show that CNN-Sig algorithm could be a good candidate in the high-dimensional regression problem where plain signature is not applicable. But since CNN-Sig will add non-linearity here, we are not able to price this option in the same way as Arribas (2018). This will be left as our future research.

Experiments - Classification

We apply the CNN-Sig model to different high-dimensional times series from Baydogan (2015) and Ruiz et al. (2021). As suggested in Ruiz et al. (2021), all experiments are compared with a benchmark model ROCKET (Dempster et al. 2020). The results are evaluated over 5 independent trials and listed in Table 7. ROCKET is known to be a fast and accurate classification method, the experiment results show that the CNN-Sig model is competitive and fast after a model selection procedure via k -fold cross validation.¹

4.4 Sentiment Analysis by Signature

In Natural Language Processing (NLP), text sentence can be regarded as sequential data. A conventional way to represent words is using high-dimensional vector, which is called word embedding. These kinds of word embedding are usually of 50, 100, 300 dimension. Using plain signature transform becomes extremely difficult because of these high dimensions. We apply our CNN-Sig model to address this problem. The dataset we use is IMDB movie reviews, (Maas et al. 2011).

This IMDB dataset contains 50,000 movie reviews, each of them is labelled by either "pos" or "neg", which represent **Positive** for **Negative** respectively. The IMDB dataset is split into training and testing evenly. For training part, we use 17,500 samples for training the model, and use the other 7500 samples as validation dataset. A 100-dimension word embedding GloVe 100d (Pennington et al. 2014) is used as the

¹ All experiments are trained on a server with Intel Core i9-9820X (3.30GHz) and four RTX 2080 Ti GPUs




Fig. 4 QQ plot for training and testing result for regression task with $d = 50$ using CNN-Sig model

Table 7 Testing accuracy, standard deviation and total training time (s) for all high-dimensional time series datasets

Datasets	ROCKET	CNN-Sig
PEMS-SF	0.810 (0.014)	0.817 (0.010)
JapaneseVowels	0.960 (0.002)	0.940 (0.017)
FingerMovement	0.500 (0.01)	0.514 (0.034)
FaceDetection	0.597 (0.004)	0.553 (0.001)
PhonemeSpectra	0.035 (0.002)	0.152 (0.006)
MotorImagery	0.620 (0.007)	0.524 (0.05)
Heartbeat	0.729 (0.011)	0.723 (0.017)
Training Time	353.5	209.1

Bold content implies the result is better compare to the other model

initial embedding, this high dimension restricts us to use plain signature transform. In our model, by setting γ to be small, we use 1 convolutional 2d layer to reduce the dimension from 100 to c paths with each of $\gamma + 1$ dimensional augmented by extra time dimension. The architecture is shown in Fig. 5.




Fig. 5 Convolutional Signature neural network model for IMDB dataset

Table 8 Testing accuracy, GPU memory usage(Gb) during training and total training time(s) on IMDB dataset

	Bi-LSTM	CNN-Sig
Accuracy	0.846 (0.013)	0.869 (0.002)
Memory	6.8	1.3
Time	401.5	292.5

Bold content implies the result is better compare to the other model

The result is shown in Table 8 and the testing accuracy has been improved to **86.9%** which is higher than the result in Toth and Oberhauser (2019) (83%) and Bidirectional LSTM (Bi-LSTM) with 2 hidden layers (0.846%). Moreover, CNN-Sig is a more efficient structure compare to Bi-LSTM in terms of training time and GPU memory usage.

We believe that the CNN-Sig model is a good candidate for feature mapping and easy to be embraced into more complex models. By applying more complicated structure, such as using attention model for Φ and a sliding window, e.g., see Morrill et al. (2020), for calculating a sequential signature transform, the accuracy can be improved.

5 Conclusion

Using the signature to summarize sequential data has been proved to be very efficient in the low-dimensional cases. However, signature transform suffers from exponential growth of the number of features with respect to the path dimension. This makes both regression and classification problem impossible in practice.

In this paper, we proposed the Convolutional Signature (CNN-Sig) model to address this problem. Using a convolutional layer, we achieve a linear growth of the number of features and preserve all information simultaneously. The experiments show that this model can be a good candidate for classifying multi-dimension sequential data. Moreover, signature has been proved experimentally to be insensitive to missing values, this property may be useful in many natural language processing (NLP) tasks. The CNN-Sig model mitigates the high-dimension problem and provides a possible way to apply the signature transforms.

Acknowledgements We thank Professor Terry Lyons for his suggestions on mathematical aspects of CNN and for sharing the recent manuscript (Salvi et al. 2021) on the kernel method with us.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Arribas, I.P. (2018). Derivatives pricing using signature payoffs Preprint is available at [arxiv: abs/1809.09466](https://arxiv.org/abs/1809.09466)
- Baydogan, M. (2015). Multivariate Time Series Classification Datasets . Available at <http://mustafabaydogan.com>, [Accessed: 2020-07-12]
- Boedihardjo, H., Geng, X., Lyons, T., & Yang, D. (2016). The signature of a rough path: uniqueness. *Advances in Mathematics*, 293, 720–737.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Chevyrrev, I., & Lyons, T. (2016). Characteristic functions of measures on geometric rough paths. *Annals of Probability*, 44(6), 4049–4082.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H. & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation
- Cho, K., van Merriënboer, B., Gulcehre, C., Bougares, F., Schwenk, H., & Bengio, Y. (2014) Learning phrase representations using rnn encoder-decoder for statistical machine translation. (p. (2014)) [arxiv: abs/1406.1078](https://arxiv.org/abs/1406.1078).
- Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics Of Control, Signals And Systems*, 2, 303–314. <https://doi.org/10.1007/BF02551274>.
- Dempster, A., Petitjean, F., & Webb, G. I. (2020). Rocket: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery*, 34(5), 1454–1495.
- Detering, N., Fouque, J. P., & Ichiba, T. (2021). Directed chain stochastic differential equations. *Stochastic Processes and their Applications*, 130, 2519–2551.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50, 987–1007.
- Funahashi, K. I. (1989). On the approximate realization of continuous mappings by neural networks. *Neural Networks*, 2(3), 183–192.
- Gyurkó, L.G., Lyons, T., Kontkowski, M., & Field, J (2013). Extracting information from the signature of a financial data stream [arxiv: abs/1307.7244](https://arxiv.org/abs/1307.7244)
- van Handel, R. (2016). *Probability in High Dimension. APC 550 Lecture Notes* (p. Princeton University). <https://web.math.princeton.edu/~rvan/APC550.pdf>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.*, 9(8), 1735–1780.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), 13–30.
- Kidger, P., Bonnier, P., Perez Arribas, I., Salvi, C. & Lyons, T. (2019). Deep signature transforms. In: *Advances in Neural Information Processing Systems* 32, pp. 3105–3115. Curran Associates, Inc. . <http://papers.nips.cc/paper/8574-deep-signature-transforms.pdf>
- Kidger, P., & Lyons, T. ((2020,7,9),). Universal Approximation with Deep Narrow Networks. *Proceedings Of Thirty Third Conference On Learning Theory.*, 125, 2306–2327. <https://proceedings.mlr.press/v125/kidger20a.html>.
- Kidger, P., & Lyons, T. (2020). Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU. <https://github.com/patrick-kidger/signatory>
- Kiraly, F. J., & Oberhauser, H. (2019). Kernels for sequentially ordered data. *Journal of Machine Learning Research*, 20(31), 1–45.
- Levin, D.A., Lyons, T. & Ni, H.(2013). Learning from the past, predicting the statistics for the future, learning an evolving system Preprint is available at [arxiv: abs/1309.0260](https://arxiv.org/abs/1309.0260)
- Liao, S., Lyons, T., Yang, W. & Ni, H.(2019). Learning stochastic differential equations using RNN with log signature features Preprint is available at [arxiv: abs/1908.08286](https://arxiv.org/abs/1908.08286)
- Lyons, T., Nejad, S., & Arribas, I. P. (2019). Numerical method for model-free pricing of exotic derivatives using rough path signatures. *Applied Mathematical Finance*, 26, 583–597.
- Lyons, T., Nejad, S., & Arribas, I. P. (2020). Non-parametric pricing and hedging of exotic derivatives. *Applied Mathematical Finance*, 27, 457–494.
- Lyons, T., Ni, H., & Oberhauser, H. (2014). *A feature set for streams and an application to high-frequency financial tick data*. New York: Association for Computing Machinery. <https://doi.org/10.1145/2640087.2644157>.

- Lyons, T., & Qian, Z. (2002). *System Control and Rough Paths*. Oxford mathematical monographs: Clarendon Press.
- Lyons, T. J., Caruana, M., & Lévy, T. (2007). Differential equations driven by rough paths. *Lecture Notes in Mathematics Differential Equations Driven by Rough Paths: École d'Été de Probabilités de Saint-Flour XXXIV - 2004* (Vol. 1908, pp. 81–93). Berlin Heidelberg: Springer.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., & Potts, C. (2011). *Learning word vectors for sentiment analysis* (pp. 142–150). Portland, Oregon, USA: Association for Computational Linguistics. <http://www.aclweb.org/anthology/P11-1015>.
- Min, M., & Hu, R. (2021) Signed Deep Fictitious Play for Mean Field Games with Common Noise. *Proceedings Of The 38th International Conference On Machine Learning*. **139** pp. 7736-7747 (2021,7,18), <http://proceedings.mlr.press/v139/min21a.html>
- Morrill, J., Fermanian, A., Kidger, P. & Lyons, T. (2020) A generalized signature method for time series Preprint is available at [arxiv: abs/2006.00873](https://arxiv.org/abs/2006.00873)
- Pennington, J., Socher, R. & Manning, C.D. (2014). Glove: Global vectors for word representation. In: *Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 <http://www.aclweb.org/anthology/D14-1162>
- Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2018). Numerical Gaussian processes for time-dependent and non-linear partial differential equations. *SIAM Journal on Scientific Computing*, *40*, A172–A198.
- Ruiz, A. P., Flynn, M., Large, J., Middlehurst, M., & Bagnall, A. (2021). The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery*, *35*(2), 401–449.
- Salvi, C., Cass, T., Foster, J., Lyons, T., & Yang, W. (2021). The Signature Kernel Is the Solution of a Goursat PDE. *SIAM J. Math. Data Sci.*, *3*(3), 873–899.
- Toth, C. & Oberhauser, H. (2019). Bayesian learning from sequential data using Gaussian processes with signature covariances (2019) Preprint is available at [arxiv: abs/1906.08215](https://arxiv.org/abs/1906.08215)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L. & Polosukhin, I. (2017) Attention is all you need. In: I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds.) *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc. (2017). <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- Williams, C. K., & Rasmussen, C. E. (2006). *Gaussian processes for machine learning*. MA: MIT press Cambridge.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.