Neural-Signature Methods for Structured EHR Prediction

Andre Vauvelle (✉ andre.vauvelle.19@ucl.ac.uk)

University College London

Paidi Creed

BenevolentAI (United Kingdom)

Spiros Denaxas

University College London

Research Article

Keywords: machine learning, electronic healthcare records, signature methods

Posted Date: February 15th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1292041/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

RESEARCH

Neural-Signature Methods for Structured EHR prediction

Andre Vauvelle^{1*}, Paidi Creed² and Spiros Denaxas¹

*Correspondence:

andre.vauvelle.19@ucl.ac.uk

¹Institute of Health Informatics,
University College London, 222
Euston Road, London, UK
Full list of author information is
available at the end of the article

Abstract

Models that can effectively represent structured Electronic Healthcare Records (EHR) are central to an increasing range of applications in healthcare. Due to the sequential nature of health data, Recurrent Neural Networks (RNNs) have emerged as the dominant component within state-of-the-art architectures. The signature transform represents an alternative modelling paradigm for sequential data. This transform provides a non-learned approach to creating a fixed vector representation of temporal features and has shown strong performances across an increasing number of domains, including medical data. However, the signature method has not yet been applied to structured EHR data. To this end, we follow recent work that enables the signature to be used as a differentiable layer within a neural architecture enabling application in high dimensional domains where calculation would have previously been intractable. Using a heart failure (HF) prediction task as an exemplar, we provide an empirical evaluation of different variations of the signature method and compare against state-of-the-art baselines. This first application of neural-signature methods in real-world healthcare data shows a competitive performance when compared to strong baselines and thus warrants further investigation within the health domain.

Keywords: machine learning; electronic healthcare records; signature methods

1 Introduction

Prediction tasks defined on structured EHR data are a key focus for applications of Machine Learning in Healthcare, with the potential to improve patient outcomes through faster and more accurate diagnoses. Due to the rapidly increasing quantity and availability of EHR data, methods in deep learning are increasingly being utilised to model the complex interactions in a range of healthcare related predictive tasks. Due to the sequential nature of EHR data, RNNs have emerged as a key component in many recent state of the art methods.

The signature transform maps a path (for example a time series) onto an infinite sequence of summary statistics. It is known that these terms completely characterise the path (up to translation) and that any function on the path can be modelled arbitrarily well by a linear function on the signature [21, 32]. In a machine learning context, this makes the signature a useful feature set with which to learn from. The signature has been successful across a range of predictive tasks involving time series data [35] in particular in the medical domain [34]. However, the signature method has not yet been applied to structure EHR data, most likely due to its high dimensionality posing computational challenges. To this end, we follow recent work that enables the signature to be used as a differentiable layer within a neural

architecture enabling application in high dimensional domains where calculation would have previously been intractable [8].

In this paper we perform an empirical evaluation of signature methods as a novel alternative to RNNs for disease prediction using EHR data. We create a 90 day HF prediction task with data from the UK Biobank [43] to compare neural-signature models with various augmentations against RNN and bag of words baselines. The key results can be summarised as follows:

- Neural-signature methods are able to produce a competitive predictive performance when compared to RNN models, returning results over two separate corpora and metrics within one standard deviation
- Log-signature and lead-lag variants improve results from those similar to basic bag of words models to those comparable with RNNs
- Adding time-augmentations does not significantly effect model performance for both neural-signature and RNN models

2 Related Work

The methods previously used to address temporality in EHR can be roughly separated into three main areas;

Discretization: This consists of splitting the continuous-time variables into discrete bins. Features are then calculated from the sub-sequences within each time period. For categorical data, the most common approach is to count the number of events.

Neural approaches: Neural network approaches attempt to automatically learn a feature set that best describes the underlying data for a specific prediction task. Pham et al. [37] and Choi et al. [12] applied RNN variants to find results that reported improved performance over existing state-of-the-art methods.

RNN variants continue to play a role in more recent papers [13, 15, 19, 33, 38, 39, 45, 48]. Modifications include using bidirectional RNNs to reduce steps between dependencies, attention mechanisms to improve interpretability, facilitate combining with Convolutional Neural Network (CNN) models, and improved embeddings for visits with graph-based attention models. In all such papers, RNNs are used to handle the sequential aspect of structured EHR.

While it is clear that RNNs perform comparatively well in deep learning applications, an alternative set of methods is also worth discussing.

Sequential feature extraction: This encompasses methods that are able to extract flat features from sequential data while retaining information relating to ordering. Despite being more popular in higher frequency data modalities such as streams data, previous works using structure EHR have explored; shapelets [49] and symbolic aggregate approximation [6] for adverse drug reaction prediction. More broadly, this includes methods such as the discrete Fourier transformation [1].

It is this category that signature methods can be considered to belong. A key advantage of signature methods is a strong theoretical groundwork showing the signatures usefulness in non-parametric hypothesis testing [11] and algebraic geometry [36]. Machine learning applications have also been demonstrated in a growing variety of domains [10] including: healthcare [4, 26, 27, 34], finance [3, 23], action recognition [30, 47] and hand-writing recognition [46].

Table 1: A summary of the cohort used for HF prediction experiments.

		HF	Controls	Total
# patients		5,722	51,498	57,220
PRIMDX	# events	48,779	188,893	237,672
	Average # events per patient	8.52	3.67	4.15
	Unique tokens	-	-	2147
PRIMDX-SECDX-PROC	# events	184,036	616,066	800,102
	Average # events per patient	32.16	11.96	13.98
	Unique tokens	-	-	6235

3 Cohort

The UK Biobank [43] is a national population-based study comprising of 502,629 individuals. We extract a retrospective heart failure (HF) cohort using the same methodology as Denaxas et al. [18] which uses the previously-validated phenotyping algorithm in the CALIBER resource [17].

To form the sequential input data required for our predictive model, we extract primary and secondary diagnosis terms (ICD10), procedure terms (OPCS4) and, timestamps ("epistart") from the UK Biobank inpatient dataset. Patient events are extracted with a buffer period of 90 days before HF diagnosis (for controls this is the HF diagnosis of its matched case) to exclude highly correlated events such as end of life care [28]. Events that occur at the same time are randomly ordered.

We create two separate corpora for each patient: PRIMDX is a corpus that only contains primary diagnosis terms, and PRIMDX-SECDX-PROC also includes secondary diagnoses and procedure terms. Since the number of events in each sequence is greater for the PRIMDX-SECDX-PROC cohort, this allows us to compare each of our methods' ability to handle longer sequences with more complex and redundant information.

In Table 1 we provide a breakdown of the demographics of the matched cohort used in this study. In Appendix A we provide further details on the HF cohort extracted and the tokenization process of healthcare terms.

4 Methods

Let each patient record be denoted by the path $\mathbf{x} = (x_t^1, \dots, x_t^d)$, where each value x_t^i is real-valued and parameterised by $t \in [0, T]$.

Our objective is to classify each sequence with a binary variable which indicates whether the patient will develop heart failure within 90 days. The dimension of the path, d , is determined by the maximum number of unique tokens as we represent each token with a one-hot-vector, such that only the dimension corresponding with the index of the vocabulary is one and with zeros everywhere else.

4.1 Signature methods

The definition of the signature transform is as follows.

Let $T > 0$ and $0 < t_1 < t_2 < \dots < t_{n-1} < t_n = T$. Let $f_x = (f_x^1, \dots, f_x^d) : [0, T] \rightarrow \mathbb{R}^d$ be the unique continuous function such that $f_x(t_i) = x_i$ and is affine on the intervals between them. The signature is the infinite collection of iterated

integrals

$$\text{Sig}(\mathbf{x}) = \left(\left(\int_{0 < t_1 < \dots < t_k < T} \prod_{j=1}^k \frac{df^{i_j}}{dt}(t_j) dt_j \right)_{1 \leq i_1, \dots, i_k \leq d} \right)_{1 \leq k} . \quad (1)$$

The form of the signature in Equation 1 can be broken down to help give the reader a better understanding as done in [10]. We can start by simplifying to a single index $i \in \{1, \dots, d\}$. This reduces Equation 1 to

$$\text{Sig}(\mathbf{x})^i = \int_{0 < t < T} \frac{df^i}{dt}(t) dt. \quad (2)$$

As this is a single integral and f is affine, the equation simply resolves to the increment of i -th coordinate of the path

$$\text{Sig}(\mathbf{x})^i = x_T^i - x_0^i. \quad (3)$$

The double-iterated integral considers any pair of coordinates $i, j \in \{1, \dots, d\}$ such that

$$\text{Sig}(\mathbf{x})^{i,j} = \int_{0 < t < T} \text{Sig}(\mathbf{x})^i \frac{df^j}{dt}(t) dt \quad (4)$$

$$= \int_{0 < t_1 < t_2 < T} \frac{df^i}{dt_1}(t_1) dt_1 \frac{df^j}{dt_2}(t_2) dt_2 \quad (5)$$

where we have used Equation 2 and replaced t to denote the integration limits as

$$a < t_1 < t_2 < T = \begin{cases} 0 < t_1 < t_2 \\ 0 < t_2 < T \end{cases} \quad (6)$$

Notice that the integration limits in Equation 6 correspond to the integration over a triangle. Going further this process can continue recursively and be interpreted as integrating over an increasingly high dimensional simplex. This real number is known as the k -fold iterated integral seen in Equation 1

$$\text{Sig}(\mathbf{x})^{i_1, \dots, i_k} = \int_{0 < t < T} \text{Sig}(\mathbf{x})^{i_1, \dots, i_{k-1}} \frac{df^{i_k}}{dt}(t) dt \quad (7)$$

$$= \int_{0 < t_1 < \dots < t_k < T} \prod_{j=1}^k \frac{df^{i_j}}{dt}(t_j) dt_j \quad (8)$$

where the superscripts are members of the set

$$1, \dots, k \in \{1, \dots, d\}. \quad (9)$$

We can further simplify this form and remove the need for the integral when we consider the path as a series of linear segments in a piecewise linear path. For a single segment the signature can be expressed by the product of its increment

$$\text{Sig}(\mathbf{x})_{t,t+1}^{i_1,\dots,i_k} = \frac{1}{k!} \prod_{j=1}^k (x_{t+1}^{i_j} - x_t^{i_j}). \quad (10)$$

To calculate each signature term of the full path, we can use Chen's Identity, which states that the signature of the entire path can be calculated from the signatures of its segments [9]

$$\text{Sig}(\mathbf{x})_{0,T}^{i_1,\dots,i_k} = \sum_{j=0}^k \text{Sig}(\mathbf{x})_{0,t}^{i_1,\dots,i_j} \text{Sig}(\mathbf{x})_{t,T}^{i_{j+1},\dots,i_k}. \quad (11)$$

Using the signature as an infinite series in a machine learning pipeline would not be tractable. Instead, it is common to truncate the series to the k -th level, this is also known as the depth of the signature. This results in the finite collection of terms $\text{Sig}(\mathbf{x})^{i_1,\dots,i_k}$ where the multi-index is restricted to length N . For example a signature of depth 1 is the collection of d real numbers $\text{Sig}(\mathbf{x})^1, \dots, \text{Sig}(\mathbf{x})^d$ and a signature of depth 2 is the collection of $d + d^2$ real numbers $\text{Sig}(\mathbf{x})^1, \dots, \text{Sig}(\mathbf{x})^d, \text{Sig}(\mathbf{x})^{1,1}, \dots, \text{Sig}(\mathbf{x})^{d,d}$.

The number of terms τ , for any truncated signature of order N of a d -dimensional path, where $d \geq 1$, is the geometric series:

$$\tau = \sum_{k=0}^N d^k = \frac{d^{N+1} - 1}{d - 1} \quad (12)$$

For structured EHR data with hundreds or thousands of unique terms, this poses a significant computational issue. In the next section, we highlight a number of variations that can be used to encourage information into lower order signature terms.

In Appendix B, we provide a further breakdown of the definitions provided here and explore an example in toy data to show how the signature terms describe sequential data. Theoretically, the signature terms are proven to uniquely describe any path up to translations (Proposition 1) and act as a universal non-linearity (Proposition 2). This latter property is shared with neural networks and allows us to reduce potentially complicated non-linear relationships between variables into linear ones.

4.2 Signature variations

There is a body of variations on the standard signature transform that have been developed. Each can tailor the properties of the signature to be more suited for a certain task. Morrill et al. [35] provides an overview of possible variations of the signature together with an empirical evaluation on streams data. Given the substantially greater dimensionality of structured EHR data, we restrict our investigation to the augmentations in Table 3 and the log-signature.

4.2.1 Augmentations

An augmentation considers transforming our sequence of patient events $\mathbf{x} \in \mathbb{R}^d$ into one or several new sequences, p , whilst potentially changing the dimensionality of each path to a . In general, this can be described by the map

$$\phi : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}(\mathbb{R}^a)^p. \quad (13)$$

Table 2: Summary of augmentations.

Augmentation	a	p	Length
Time	$d + 1$	1	n
Basepoint	d	1	$n + 1$
Lead-lag	$2d$	1	$2n$
Learnt projections	a	p	n

The time augmentation consists of the concatenation of an extra dimension. As shown in Proposition 1, this can be used in the absence of any actual timestamps by simply using the index of the event in the sequence. In both cases, this removes the property of time-parameterisation invariance of the signature [29]. We also investigate applying actual time differences from prediction date to account for the irregularly sampled nature of the data. We follow Rajkomar et al. [39] by applying the parameterised scaling function, $f(\Delta T) = T_{scale} \log(\Delta T)$ capped a maximum T_{max} . T_{scale} and T_{max} control extreme time-deltas and are optimised as hyperparameters.

The basepoint [24] is used to remove the property of translational invariance. This property means that the signature of two paths separated by a constant translation will be the same. The basepoint also has a significant advantage for our pipeline as $\sim 20\%$ of pathways in the dataset used in this study have only a single event. Basepoint introduces an origin point at the start of each path and thus ensures each path has at least two points which is a requirement for calculating the signature.

The lead-lag augmentation [10, 20] adds shifted copies of the path as new coordinates. This augmentation explicitly captures the quadratic variation of the underlying process, an important concept for our data where the co-variance between medical concepts is known to be highly important to the underlying pathology of disease [16, 38]. A lag of a single timestep is described by the following augmentation

$$\phi(\mathbf{x}) = ((x_1, x_1), (x_2, x_1), (x_2, x_2), \dots, (x_t, x_t)) \in \mathcal{S}(\mathbb{R}^{2d}). \quad (14)$$

The learnt projection can be described by the affine transformation or embedding, $\theta_A \in \mathbb{R}^{a \times d}$, such that

$$\phi(\mathbf{x}) = (\theta_A x_1, \theta_A x_2, \dots, \theta_A x_n) \in \mathcal{S}(\mathbb{R}^a). \quad (15)$$

This reduces the dimensionality of the path to make the calculation of the signature tractable.

4.2.2 The log-signature

The log-signature corresponds with taking the formal logarithm of the signature in the algebra of formal power series [10]. Both the signature and its logarithm uniquely define a path (Proposition 1) but the log-signature does not hold the same universality property (Proposition 2) [30]. The log-signature maps to a smaller number of terms at each truncation level determined by Witt’s formula, which is shown in Appendix B.3.2 along with an example.

4.2.3 Deep Signatures

For the affine transformations discussed in Equation 15, we briefly described a learning process. As detailed in recent works from Bonnier et al. [8], it is possible to train the affine transformations together with the signature transform through an end-to-end neural network architecture. Here, the signature acts as a non-parametric pooling function able to extract provably useful information from sequential data.

It is possible to calculate the gradient needed in this method as the signature can be formulated as a calculation tree of differentiable operations [24, 41].

The generalised function of the neural-signature model used in this work can be written as

$$f(\mathbf{x}; \Theta) = \sigma((g^{\theta_{fc}} \circ \text{Sig}^N \circ \phi^{\theta_A})(\mathbf{x})) \quad (16)$$

where we have denoted that the learnt parameters as the weights of a fully connected neural network classifier θ_{fc} and elements of the affine transformation augmentation θ_A . The sigmoid function is used to map the output activation to a $[0,1]$ score.

5 Experiments

As baselines, we consider a bag of words model with logistic regression as a commonly used most basic model, along with a GRU model, which is comparable with the state of the art *RETAIN* [5, 13]. We also include a GRU variation that incorporates the time delta augmentation.

Additionally, we consider the following signature models: the standard signature (S) provides the baseline for further variations, the log-signature (LS) removes the universality property (the fully connected neural network classifier still guarantees this overall) but greatly reduces the number of signature terms, the lead-lag (LL) augmentation encourages information about the quadratic variation into lower-order signature terms, the add time index augmentation (ATI) provides sensitivity to parameterization, the time delta (ATD) version goes further to account for non-uniform sampling rates. We limited the exploration on augmentations to the above after initial testing on validation data found the lead-lag augmentation to be most influential.

We use two metrics for evaluation; area under the receiver operator curve (AUROC) and area under the precision-recall curve (AUPRC).

Previous studies have shown that the AUROC can provide misleading results when there is considerable data imbalance, mainly if the number of negative examples is high, and we have a preference for identifying true positive examples [42]. This issue exists in our task due to the 1:9 case-control split and the increased benefit

Table 3: Sequential models compared.

Variants	Description
S^2	Standard Signature with depth 2
LS^2	Log-signature - a condensed form of the signature
$LS^2 + LL$	LS with lead-lag augmentation - extracts quadratic variation
$LS^2 + LL + ATI$	$LS + LL$ with time index - removes time reparameterisation invariance
$LS^2 + LL + ATD$	$LS + LL$ with time delta - accounts for non-uniform sampling rate
$LS^3 + LL + ATD$	Increased truncation depth - more complex, sequential features
BoW OH LR	Most basic model with one-hot bagging and logistic regression
GRU	Baseline gated recurrent unit sequential model
GRU + ATD	Takes into account time differences

of correctly identifying HF cases over correctly identifying controls. The result of this class imbalance can cause AUROC to become inflated due to a high number of true negative cases. AUPRC is an alternative metric that captures the trade-off between precision and recall. Crucially, it ignores the number of true negatives allowing changes in performance to be seen without being diluted as in AUROC.

The signature variations explored are summarised together with the baselines in Table 3. Common to each model is the architecture shown in Figure 1. Further details on implementation, including initialisation, activation functions, optimisation, hyperparameters, regularisation, and other such related details are found in Appendix C.

6 Results

Table 4: Best performing models maximised for AUROC during hyperparameter optimisation. Performance metrics evaluated on the test dataset with ± 1 standard deviation calculated from cross validation scores.

Model	PRIMDX AUROC	AUPRC	# Params	PRIMDX-SECDX-PROC AUROC	AUPRC	# Params
S^2	0.682 ± 0.011	0.211 ± 0.015	525,296	0.695 ± 0.005	0.247 ± 0.003	258,098
LS^2	0.687 ± 0.005	0.248 ± 0.014	637,906	0.704 ± 0.011	0.245 ± 0.019	167,101
$LS^2 + LL$	0.703 ± 0.007	0.241 ± 0.005	216,404	0.728 ± 0.011	0.300 ± 0.014	592,528
$LS^2 + LL + ATI$	0.700 ± 0.009	0.249 ± 0.018	864,817	0.724 ± 0.007	0.303 ± 0.009	545,474
$LS^2 + LL + ATD$	0.695 ± 0.013	0.254 ± 0.014	596,364	0.730 ± 0.008	0.304 ± 0.009	1,504,205
$LS^3 + LL + ATD$	0.698 ± 0.007	0.244 ± 0.017	7,015,643	0.734 ± 0.008	0.300 ± 0.009	4,037,181
BoW OH LR	0.699 ± 0.007	0.252 ± 0.006	-	0.690 ± 0.015	0.300 ± 0.015	-
GRU	0.701 ± 0.009	0.241 ± 0.007	795,390	0.731 ± 0.008	0.293 ± 0.017	415,892
GRU + ATD	0.710 ± 0.010	0.275 ± 0.012	668,997	0.733 ± 0.008	0.309 ± 0.012	489,151

From Table 4, we observe similar predictive performance across signature models using lead-lag augmentations and GRU models over all corpora, with all metrics from the two sets of models remaining within one standard deviation of performance

seen on the validation data. All models perform the same or better on the larger PRIMDX-SECDX-PROC cohort, but more complex models gain a more significant benefit from the added data.

The addition of time augmentations does not show a consistent improvement in performance over just applying the lead-lag augmentation, and there is no consistent difference between adding a time index and time delta. Increasing the depth of the signature to three also shows no significant increase in performance. Signature models perform similar to the bag of words baselines without the lead-lag transform.

6.1 Data ablation study

Our final set of experiments evaluates how the models perform as the volume of data is reduced. For the data ablation study, each trial randomly samples a proportion of the training and validation dataset. For each proportion a new set of hyperparameters is found for each model.

The model parameters and hyperparameters are trained using 5-fold cross validation on the sub-sample while the remaining data is unused. The ablation study test data remains the same as the main experiment.

Again, results are broadly similar for both models except for three points where the two models produce performance outside of one standard deviation of validation performance. Notably, at 20% data ablation for AUPRC, the signature model has a 21.0% higher score with 0.283 versus 0.237 for the GRU. Overall, both models' performance begins to saturate at $\sim 20\%$ for both metrics, and the results show no conclusive trend as to which model performs best as the amount of training data is reduced.

7 Discussion

Given the properties that the two methods share these results might not come as a surprise. However, without the lead-lag augmentation the performance of the signature models drop significantly. This confirms the prior belief that the quadratic variance of the path plays an important role in structured EHR HF prediction. This could correspond with encouraging features that describe changing comorbidities to be present in the lower order terms of the signature.

For the data ablation study, we expected the signature model to outperform the GRU baseline however, results for both methods are similar. Our prior hypothesis was partly motivated by the success of signature methods in previous shallow machine learning tasks [34]. A key difference in our task could be the high dimensionality and reliance on embeddings to make the signature tractable. The need to train these embeddings is likely data-intensive but could benefit from initialisation using pretrained *word2vec* embeddings as has been shown for RNNs [12].

7.1 Comparison to previous literature

Comparing the results in this paper directly to previous work is challenging due to the use of different underlying study designs, populations, and incomplete definitions of cohorts and outcomes [44]. We note that previous works investigating sequential models for predicting HF on structured EHR data have found greater performance [5, 13, 40]. In particular, these works also show a more significant performance difference between bag of words baselines and RNN based architectures.

Again, differences in the features and data sources used make comparisons difficult. For example, we could compare against Solares et al. Ayala Solares et al. [5], which also uses data from a multi-center UK EHR data source and achieves 0.951 AUROC using the RNN based model RETAIN. However, we must consider that the authors also include primary care and demographics data, which could influence prediction performance independent of model choice. The large US multi-center study by Rasmy et al. [40] show RETAIN achieving a more comparable AUROC of 0.769 on US healthcare care with a balanced cohort with 14,500 cases and with only diagnosis codes provided for as prediction input. The same model achieved an AUROC of 0.822 when trained and tested on the full cohort of 152,790 cases and 1,152,517 controls with diagnoses, demographic, medication, and surgery data.

8 Conclusion

Given the prevalence of RNNs in current structured EHR architectures, any improvement in this fundamental component is likely to influence future work significantly. A substantial body of theory motivates the use of signature transforms to represent sequential data, and previous works have shown them to have strong empirical performance. In particular, recent works on neural-signature architectures have enabled their applications on high-dimensional data.

This work is the first to show that neural-signature methods with dimensional reduction before the transform are competitive on high dimensional structured EHR data. Using an HF prediction task, we evaluated the signature transforms as an alternative to RNNs that provide a predictive and compact representation of sequential structured EHR data. We show that the signature achieves comparable performance to RNNs and that the performance of both models saturates with a similar number of training examples. While the signature originates from perhaps abstract theory, empirically, it can successfully compete with the current state-of-the-art architectures.

Appendix A: Cohort

The HF cohort excludes patients with: self-reported prevalent HF cases, those who died during the study period, heart failure diagnoses outside the study period between 1997 and 2015. The cases are matched to controls on assessment center, year of recruitment, sex, and year of birth. Controls use an index date that is assigned to the date of HF diagnosis for its matched case. We randomly sampled the total potential control population of 496,892 with the above matching criteria to create a cohort with a ratio of 1:9 cases to controls. This resulted in 5,722 cases and 51,498 controls, where 15 cases could not be matched.

A.1 Tokenization

We follow a similar practice to NLP tokenization, using tags to identify whether a code comes from a particular corpus. This allows us to distinguish between primary and secondary codes with the same ICD10 term and introduce additional tokens to indicate if there is no data present. The tokenization process also applies two measures to reduce the vocabulary size; limit the length of all terms to 4 characters and a minimum count of 5. This is possible because many ICD10 and OPCS contain extension codes or additional sub-chapter codes which are not commonly used

across healthcare providers and often contain details that indicate minor differences between terms. The tokenizer vocabulary is fitted on the training data, and any subsequent code that is not present in the vocabulary is assigned an out-of-vocabulary token in place.

Appendix B: Signature Methods

B.1 A geometric intuition and exploration in toy data

Previously in Equation 3 we showed that the first level resolves to the increment of a path in a single dimension. It is also possible to find a geometric interpretation for 2nd level terms in Equation 2. For 2nd level terms where $i = j$ we find $\text{Sig}(\mathbf{x})^{i,j} = (x_T^i - x_0^i)^2/2$ which is simply the area of a triangle sided with the increment in the coordinate. If $i \neq j$ then $\text{Sig}(\mathbf{x})^{i,j} = (x_T^i - x_0^j)^2/2$. Another geometric comparison can be made here by considering the Lévy area, which is the signed area shown in Figure 2. This area illustrates the 2nd order terms through the following equation

$$A_{\text{levy}} = \frac{1}{2}(\text{Sig}(\mathbf{x})^{1,2} - \text{Sig}(\mathbf{x})^{2,1}). \quad (17)$$

The Figure 2 can be related to our task if we consider the line (path 2) and its cord (path 1) to be two separate patient pathways with three unique events (1,2,3), but for path 2 the order of events 1 and 2 are swapped. When classifying a patient pathway with an HF prediction, the order of events may hold vital information. However, when we calculate the terms of the signature in Table 5, we see that the two paths are only separable when taking a signature of depth 3.

Taking the geometric interpretation to higher orders becomes more complicated, but for additional interpretable relations, we refer the reader to [10] which expands on the ideas here to include relations to statistic moments. It is clear that these signature terms are extracting information about the order and that higher-order signature terms can be used to discriminate different paths better. It is also possible to observe how the terms of the signature change by drawing a path using an online tool^[1].

Table 5: Signature terms up the 3rd level for the paths in Figure 2.

Sig Terms	Path 1	Path 2	Different?
(1,)	3	3	FALSE
(2,)	3	3	FALSE
(1, 1)	4.5	4.5	FALSE
(1, 2)	4.5	4.5	FALSE
(2, 1)	4.5	4.5	FALSE
(2, 2)	4.5	4.5	FALSE
(1, 1, 1)	4.5	4.5	FALSE
(1, 1, 2)	4.5	4	TRUE
(1, 2, 1)	4.5	5.5	TRUE
(1, 2, 2)	4.5	5.5	TRUE
(2, 1, 1)	4.5	4	TRUE
(2, 1, 2)	4.5	2.5	TRUE
(2, 2, 1)	4.5	5.5	TRUE
(2, 2, 2)	4.5	4.5	FALSE

^[1]<http://imanolperez.pythonanywhere.com/>

Figure 2: A plot of two event paths with path 1 also being the chord that joins the two end points of path 2. The Levy area is negative above the chord and positive above.

B.2 Key properties and caveats for signature methods

A natural extension to our toy example would be to ask if the signature can be used to discriminate between any path. In fact, it has been shown that a path is essentially defined by its signature and that almost no information is lost. There are two key properties of the signature that make its use interesting for path-like data.

Proposition 1 (*Uniqueness of signature [22]*). *Given $\mathbf{x} = (x_1, \dots, x_n)$ then the map $\mathbf{x} \rightarrow \text{Sig}^\infty((1, x_1), (2, x_2), \dots, (n, x_n))$ uniquely determines $\mathbf{x}$ up to translations.* For the purpose of this statement we have introduced the two concepts of time-augmenting a path and of translational invariance. Both of which will be expanded on in Section B.3.

Proposition 2 (*Universal non-linearity [3]*) *Let F be a real-valued continuous function on continuous piecewise smooth paths in $\mathbb{R}^d$ and $\mathcal{K}$ be a compact set of such paths. Furthermore assume that $x_0 = 0$ for all $\mathbf{x} \in \mathcal{K}$. (To remove the translational invariance.) Let $\epsilon > 0$. Then there exists a linear functional L such that for all $\mathbf{x} \in \mathcal{K}$*

$$|F(\mathbf{x}) - L(\text{Sig}(\mathbf{x}))| < \epsilon. \quad (18)$$

In other words, we know that we can find a linear function that when applied to the signature of a path can approximate any function applied to that same path. This powerful property is shared with neural network approaches and allows us to reduce potentially complicated non-linear relationships between variables into the linear ones.

For streams data such as heart rate or accelerator signals, this is ordinarily low, but for categorical event data, dimensionality can often be orders of magnitude higher. Our next section goes on to outline several strategies that have been used to encourage more information into the lower order signature terms.

In this section, we have covered the signature's key properties that motivate their use for our task; however, there are several additional useful properties related to machine learning, some of which we will cover while motivating augmentations in the next section. Otherwise, we refer the interested reader to Appendix A of [8] for an expanded summary.

B.3 Variations on the signature transform

There is a growing body of variations on the standard signature transform that have been developed. Each can tailor the properties of the signature to be more suited for a certain task. Morrill et al. [35] gives an overview of the variations on the signature transform and proposes a generalised framework for extracting signature features for time-series data.

$$\mathbf{y}_{i,j,k} = (\rho_{post} \circ \text{Sig}^N \circ \rho_{pre} \circ W^{i,j} \circ \phi^k)(\mathbf{x}). \quad (19)$$

Here we have introduced a number of new terms. Each of which will be explained in more detail and discussed in relation to our data. Briefly; ϕ is known as an augmentation, W describes a windowing operation, ρ describes a re-scaling operation

both pre and post calculating the signature. Finally, Sig^N represents either the now familiar signature operation up to depth N or the log-signature variation, which will be introduced later in this section. The indices (i, j) refer to the signature terms of windowed operations, with $\mathbf{y}_{i,j,k}$ reducing to the more familiar signature terms up to level k , $\mathbf{y}_k$, when a global window is used. For our work, we will focus on only two aspects of this generalised framework, augmentations and the log-signature, leaving the exploration of the additional steps to future work.

B.3.1 Augmentation

An augmentation considers transforming our sequence of patient events $\mathbf{x} \in \mathbb{R}^d$ into one or several new sequences, p , whilst potentially changing each paths dimensionality to a . In general this can be described by the map

$$\phi : \mathcal{S}(\mathbb{R}^d) \rightarrow \mathcal{S}(\mathbb{R}^a)^p. \quad (20)$$

There are a number of different augmentations that have been used for continuous time series data.

The augmentations in Table 2 are broadly separated in two categories. Fixed augmentations consider using a transformation that does not vary after initialisation and learnt augmentations which include learnable parameter weights that are trained as part of the model fitting process. We start by explaining the two previously mentioned augmentations, time and basepoint augmentations.

The time augmentation is the addition of an extra dimension that is dependant on the order of the sequence. As shown in Proposition 1, this can be used in the absence of any actual timestamps by simply using the index of the event in the sequence

$$\phi(\mathbf{x}) = ((1, x_1), \dots, (n, x_n)) \in \mathcal{S}(\mathbb{R}^{d+1}) \quad (21)$$

or if timestamps are available

$$\phi_t(\mathbf{x}) = ((t_1, x_1), \dots, (t_n, x_n)) \in \mathcal{S}(\mathbb{R}^{d+1}). \quad (22)$$

In both cases this removes the property of time-parameterisation invariance of the signature [29]. Meaning that without time augmentation the signature only encodes the order in which events arrive and does not consider the when event arrives. This could potentially be a significant factor for our task. For example, consider two patient records of the same sequence of HF related inpatient admissions, if in one sequence the frequency of these admission was much higher it might suggest a increased disease progression. Using the actual timestamps allows the signature to account for the irregularly sampled nature of the data.

The basepoint [24] and invisibility augmentations [47] are both created with the goal of removing the property of translational invariance. This says that the signature of two paths separated by a constant translation will be the same. In [35] a comparison shows that the invisibility augmentation has essentially the same performance as the basepoint but increases the dimensionality of the path, which is of

concern since the signature scales with $\mathcal{O}(d^N)$ as seen in Equation 12. The basepoint also has a significant advantage for our pipeline as $\sim 20\%$ of pathways in the dataset used in this study have only a single event. Basepoint introduces an origin point at the start of each path and thus ensures each path has at least two points which is a requirement for calculating the signature.

$$\phi_t(\mathbf{x}) = (0, x_1, \dots, x_n) \in \mathcal{S}(\mathbb{R}^d). \quad (23)$$

The lead-lag augmentation [10, 20] adds shifted copies of the path as new coordinates. This explicitly captures the quadratic variation of the underlying process, an important concept for our data where the co-variance between medical concepts or otherwise described as comorbidities, are known to be highly important to the underlying pathology of disease [16, 38]. A lag of a single timestep is described by the following augmentation

$$\phi(\mathbf{x}) = ((x_1, x_1), (x_2, x_1), (x_2, x_2), \dots, (x_n, x_n)) \in \mathcal{S}(\mathbb{R}^{2d}). \quad (24)$$

Lead-lag comes at a high cost as the dimensionality and length of the path is doubled. The remaining augmentations provide a potential remedy for this as they propose reducing the dimensionality of the path before calculating the signature. We only consider stream-preserving neural networks that take one projection, p , as the dimensionality of our paths is orders of magnitude higher than that considered in [35].

Stream-preserving neural networks have also been used as an initial stage in previous literature on structured EHR data HF prediction tasks and provide a strong basis for comparison [5, 14, 18]. In its most basic form, a stream preserving neural network can be created by introducing a learnable affine transformation, $\theta_A \in \mathbb{R}^{a \times d}$, such that

B.3.2 The log-signature

The log-signature is a compressed version of the signature that only retains the Lyndon words of the index series [2]. As an example, we show the terms of the log signature up to depth 2 as

$$\begin{aligned} \log \text{Sig}^2(\mathbf{x}) = & (1, \text{Sig}(\mathbf{x})^1, \dots, \text{Sig}(\mathbf{x})^d, \\ & \frac{1}{2}(\text{Sig}(\mathbf{x})^{1,2} - \text{Sig}(\mathbf{x})^{2,1}), \dots, \\ & \frac{1}{2}(\text{Sig}(\mathbf{x})^{1,d} - \text{Sig}(\mathbf{x})^{d,1}), \dots, \\ & \frac{1}{2}(\text{Sig}(\mathbf{x})^{d-1,d} - \text{Sig}(\mathbf{x})^{d,d-1})) \end{aligned} \quad (25)$$

Here we can notice that the 2nd order terms of the log signature directly correspond with the Lévy area in Equation 17.

The number of terms in this more compact form of the signature can be generalised with the following.

Proposition 3 *The number of terms of the truncated log-signature τ_{log} , of order N of a d -dimensional path, where $d \geq 1$ is given by Witt’s formula [31]*

$$\tau_{log} = w(d, N) = \sum_{k=1}^N \frac{1}{k} \sum_{i|k} \mu\left(\frac{k}{i}\right) d^i \quad (26)$$

where μ is the Möbus function.

This results in almost a third of the output signature terms when calculating the signature for a path with $d = 50$ at depth 3 ($\tau_{log} = 42,925$ verses $\tau = 127,550$).

Appendix C: Experiments

C.1 Augmentations

After tokenization, we project the sequence into a lower-dimensional representation with an embedding, which is the learnable affine transformation we describe in Equation 15. For neural-signature and GRU baselines, this is initialized with a Xavier distribution and trained in an end-to-end fashion and updated during gradient descent.

We experiment with applying different combinations of previously discussed augmentations: lead-lag, index, or time delta parameterization. If both lead-lag and time parameterization are applied, then time parameterization is always done first, such that $\phi(x) \in \mathbb{R}^{2(a+1)}$. This sequence of embedded codes is used as input to the encoder stage.

C.2 Classifier

Dropout is applied to the outputs of fully connected hidden layers and determined by hyperparameter optimization.

The bag of words models use a logistic regression classifier with the sklearn implementation of L-BFGS-B and L2 regularisation.

C.3 Implementation

The experiments were implemented using: PyTorch and Signatory [24]. Models were trained on various GPU cards available through the UCL Department of Computer Science High Performance Computing Cluster, using distributed asynchronous hyperparameter optimization and MongoDB. Code used for the project, including full details on hyperparameter bounds are available at <https://github.com/andre-vauvelle/doctor-signature>.

C.4 Training and validation schema

All models are trained by minimizing binary cross-entropy loss with Adam [25]. We also stop training after five epochs if validation loss does not improve. The best model is then taken for evaluation on the test dataset.

Nested cross validation is used to evaluate our models and baselines. The data is split into training, validation, and held-out test sets with a 90% training and validation, 10% test data split. Training and validation sets are used in an inner loop with 5-fold cross validation to train model hyperparameters. Each fold is stratified such

that we preserve the percentage of samples for each class. We use a Bayesian hyper-optimization algorithm, Tree-structured Parzen Estimator, to reduce the number of evaluations needed. We fix the number of trails for each model at 100 [7]. We then select the hyperparameters with the highest average AUROC score across a 5-fold. Test data is unseen until the final evaluation, where we use it to make predictions using models trained with optimal hyperparameters from the inner validation loop on one fold of data.

Declarations

Ethics approval and consent to participate

Ethical approval for this study was provided from the UK Biobank Access Review Board, reference 58356 "Defining and redefining human disease at scale: an atlas of the human phenotype". All analyses have been performed in accordance with the Declaration of Helsinki.

Consent for publication

All UK Biobank participants have consented to the study. More information is available here: <https://www.ukbiobank.ac.uk/consent>.

Availability of data and material

The data that support the findings of this study are available from the UK Bioank (<https://www.ukbiobank.ac.uk/>) but restrictions apply to the availability of these data, which were used under license for the current study, and so are not publicly available. For re-using these data, an application must be made directly to the UK Biobank. Code is available at <https://github.com/andre-vauvelle/doctor-signature>.

Competing interests

SD has received consultation fees from BAI which were not related to the work in this paper. PC was a BAI employee at the time this research was undertaken. AV has PhD studentship funded by BAI. The authors declare that they have no other competing interests.

Funding

AV is supported by a Benevolent AI studentship. SD is funded from: a) Health Data Research UK, which receives its funding from HDR UK Ltd (HDR-9006) funded by the UK Medical Research Council, Engineering and Physical Sciences Research Council, Economic and Social Research Council, Department of Health and Social Care (England), Chief Scientist Office of the Scottish Government Health and Social Care Directorates, Health and Social Care Research and Development Division (Welsh Government), Public Health Agency (Northern Ireland), British Heart Foundation (BHF) and the Wellcome Trust. b) the BigData@Heart Consortium, funded by the Innovative Medicines Initiative-2 Joint Undertaking under grant agreement No 116074., c) the NIHR UCLH Biomedical Research Centre, d) an Alan Turing Fellowship (EP/N510129/1), e) the Longitudinal Health and Wellbeing COVID-19 National Core Study, which was established by the UK Chief Scientific Officer in October 2020 and funded by UK Research and Innovation (grant references MC.PC.20030 and MC.PC.20059), by the Data and Connectivity National Core Study, led by Health Data Research UK in partnership with the Office for National Statistics and funded by UK Research and Innovation (grant reference MC.PC.20058), and by the CONVALESCENCE study of long COVID, which is funded by NIHR/UKRI, g) The British Heart Foundation Data Science Centre (grant No SP/19/3/34678, awarded to Health Data Research (HDR) UK).

Author's contributions

AV wrote the manuscript, conducted the analysis and interpreted findings. SD created the phenotyping algorithms and phecodes. SD and PC supervised the project and assisted with manuscript revisions. All authors have read and approved the final manuscript.

Acknowledgements

We thank Dr. Eda Ozyigit, Dr Hao Ni and Dr James Morrill for their insightful discussion during this work.

Author details

¹Institute of Health Informatics, University College London, 222 Euston Road, London, UK. ²BenevolentAI, 4-8 Maple St, London, UK.

References

1. Rakesh Agrawal, Christos Faloutsos, and Arun Swami. Efficient similarity search in sequence databases. In David B. Lomet, editor, *Foundations of Data Organization and Algorithms*, Lecture Notes in Computer Science, pages 69–84, Berlin, Heidelberg, 1993. Springer. ISBN 978-3-540-48047-1. .
2. Carlos Améndola, Peter Friz, and Bernd Sturmfels. Varieties of Signature Tensors. *Forum of Mathematics, Sigma*, 7:e10, 2019. ISSN 2050-5094. . URL <http://arxiv.org/abs/1804.08325>. arXiv: 1804.08325.
3. Imanol Perez Arribas. Derivatives pricing using signature payoffs. *arXiv:1809.09466 [q-fin]*, September 2018. URL <http://arxiv.org/abs/1809.09466>. arXiv: 1809.09466.

4. Imanol Perez Arribas, Kate Saunders, Guy Goodwin, and Terry Lyons. A signature-based machine learning model for bipolar disorder and borderline personality disorder. *Translational Psychiatry*, 8(1):274, December 2018. ISSN 2158-3188. . URL <http://arxiv.org/abs/1707.07124>. arXiv: 1707.07124.
5. Jose Roberto Ayala Solares, Francesca Elisa Diletta Raimondi, Yajie Zhu, Fatemeh Rahimian, Dexter Canoy, Jenny Tran, Ana Catarina Pinho Gomes, Amir H. Payberah, Mariagrazia Zottoli, Milad Nazarzadeh, Nathalie Conrad, Kazem Rahimi, and Gholamreza Salimi-Khorshidi. Deep learning for electronic health records: A comparative review of multiple deep neural architectures. *Journal of Biomedical Informatics*, 101:103337, January 2020. ISSN 1532-0464. . URL <http://www.sciencedirect.com/science/article/pii/S1532046419302564>.
6. Francesco Bagattini, Isak Karlsson, Jonathan Rebane, and Panagiotis Papapetrou. A classification framework for exploiting sparse multi-variate temporal features with application to adverse drug event detection in medical records. *BMC Medical Informatics and Decision Making*, 19, January 2019. ISSN 1472-6947. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6327495/>.
7. James S. Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for Hyper-Parameter Optimization. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2546–2554. Curran Associates, Inc., 2011. URL <http://papers.nips.cc/paper/4443-algorithms-for-hyper-parameter-optimization.pdf>.
8. Patric Bonnier, Patrick Kidger, Imanol Perez Arribas, Cristopher Salvi, and Terry Lyons. Deep Signature Transforms. *arXiv:1905.08494 [cs, stat]*, October 2019. URL <http://arxiv.org/abs/1905.08494>. arXiv: 1905.08494.
9. Kuo-Tsai Chen. Integration of Paths—A Faithful Representation of Paths by Noncommutative Formal Power Series. *Transactions of the American Mathematical Society*, 89(2):395–407, 1958. ISSN 0002-9947. . URL <https://www.jstor.org/stable/1993193>. Publisher: American Mathematical Society.
10. Ilya Chevyrev and Andrey Kormilitzin. A Primer on the Signature Method in Machine Learning. *arXiv:1603.03788 [cs, stat]*, March 2016. URL <http://arxiv.org/abs/1603.03788>. arXiv: 1603.03788.
11. Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *arXiv:1810.10971 [math, stat]*, October 2018. URL <http://arxiv.org/abs/1810.10971>. arXiv: 1810.10971.
12. Edward Choi, Mohammad Taha Bahadori, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. Doctor AI: Predicting Clinical Events via Recurrent Neural Networks. *arXiv:1511.05942 [cs]*, September 2016. URL <http://arxiv.org/abs/1511.05942>. arXiv: 1511.05942.
13. Edward Choi, Mohammad Taha Bahadori, Jimeng Sun, Joshua Kulas, Andy Schuetz, and Walter Stewart. RETAIN: An Interpretable Predictive Model for Healthcare using Reverse Time Attention Mechanism. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3504–3512. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6321-retain-an-interpretable-predictive-model-for-healthcare-using-reverse-time-attention-mechanism.pdf>.
14. Edward Choi, Andy Schuetz, Walter F. Stewart, and Jimeng Sun. Medical Concept Representation Learning from Electronic Health Records and its Application on Heart Failure Prediction. *arXiv:1602.03686 [cs]*, November 2016. URL <http://arxiv.org/abs/1602.03686>. arXiv: 1602.03686.
15. Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F. Stewart, and Jimeng Sun. GRAM: Graph-based Attention Model for Healthcare Representation Learning. *arXiv:1611.07012 [cs, stat]*, April 2017. URL <http://arxiv.org/abs/1611.07012>. arXiv: 1611.07012.
16. Edward Choi, Zhen Xu, Yujia Li, Michael W. Dusenberry, Gerardo Flores, Yuan Xue, and Andrew M. Dai. Learning the Graphical Structure of Electronic Health Records with Graph Convolutional Transformer. *arXiv:1906.04716 [cs, stat]*, January 2020. URL <http://arxiv.org/abs/1906.04716>. arXiv: 1906.04716.
17. Spiros Denaxas, Julie George, Emily Herrett, Anoop D. Shah, Dipak Kalra, Aroon D. Hingorani, Mika Kivimäki, Adam D. Timmis, Liam Smeeth, and Harry Hemingway. Data Resource Profile: Cardiovascular disease research using linked bespoke studies and electronic health records (CALIBER). *International Journal of Epidemiology*, 41(4):1625–1638, December 2012. ISSN 0300-5771. . URL <https://researchportal.helsinki.fi/en/publications/data-resource-profile-cardiovascular-disease-research-using-linker>. Publisher: Oxford University Press.
18. Spiros Denaxas, Pontus Stenetorp, Sebastian Riedel, Maria Pikoula, Richard Dobson, and Harry Hemingway. Application of Clinical Concept Embeddings for Heart Failure Prediction in UK EHR data. *arXiv:1811.11005 [cs, stat]*, November 2018. URL <http://arxiv.org/abs/1811.11005>. arXiv: 1811.11005.
19. Martin Ester and Dino Pedreschi, editors. *Health-ATM: A Deep Architecture for Multifaceted Patient Health Record Representation and Risk Prediction*. Society for Industrial and Applied Mathematics, Philadelphia, PA, May 2018. ISBN 978-1-61197-532-1. . URL <https://epubs.siam.org/doi/book/10.1137/1.9781611975321>.
20. Guy Flint, Ben Hambly, and Terry Lyons. Discretely sampled signals and the rough Hoff process. *Stochastic Processes and their Applications*, 126(9):2593–2614, September 2016. ISSN 0304-4149. . URL <http://www.sciencedirect.com/science/article/pii/S0304414916000429>.
21. Peter K. Friz and Nicolas B. Victoir. *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 2010. ISBN 978-0-521-87607-0. . URL <https://www.cambridge.org/core/books/multidimensional-stochastic-processes-as-rough-paths/DD9063039477D9A21F571861A5BB5D53>.
22. Ben Hambly and Terry Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, 171(1):109–167, March 2010. ISSN 0003-486X. . URL <http://arxiv.org/abs/math/0507536>. arXiv: math/0507536.
23. Jasdeep Kalsi, Terry Lyons, and Imanol Perez Arribas. Optimal execution with rough path signatures. *arXiv:1905.00728 [q-fin]*, May 2019. URL <http://arxiv.org/abs/1905.00728>. arXiv: 1905.00728.
24. Patrick Kidger and Terry Lyons. Signatory: differentiable computations of the signature and logsignature transforms, on both CPU and GPU. *arXiv:2001.00706 [cs, stat]*, January 2020. URL <http://arxiv.org/abs/2001.00706>. arXiv: 2001.00706.

25. Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, January 2017. URL <http://arxiv.org/abs/1412.6980>. arXiv: 1412.6980.
26. A. B. Kormilitzin, K. E. A. Saunders, P. J. Harrison, J. R. Geddes, and T. J. Lyons. Application of the Signature Method to Pattern Recognition in the CEQUEL Clinical Trial. *arXiv:1606.02074 [stat]*, June 2016. URL <http://arxiv.org/abs/1606.02074>. arXiv: 1606.02074.
27. Andrey Kormilitzin, Kate E. A. Saunders, Paul J. Harrison, John R. Geddes, and Terry Lyons. Detecting early signs of depressive and manic episodes in patients with bipolar disorder using the signature-based model. *arXiv:1708.01206 [stat]*, August 2017. URL <http://arxiv.org/abs/1708.01206>. arXiv: 1708.01206.
28. Konstantina Kourou, Themis P. Exarchos, Konstantinos P. Exarchos, Michalis V. Karamouzis, and Dimitrios I. Fotiadis. Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal*, 13:8–17, November 2014. ISSN 2001-0370. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4348437/>.
29. Daniel Levin, Terry Lyons, and Hao Ni. Learning from the past, predicting the statistics for the future, learning an evolving system. *arXiv:1309.0260 [q-fin]*, March 2016. URL <http://arxiv.org/abs/1309.0260>. arXiv: 1309.0260.
30. Shujian Liao, Terry Lyons, Weixin Yang, and Hao Ni. Learning stochastic differential equations using RNN with log signature features. *arXiv:1908.08286 [cs, stat]*, September 2019. URL <http://arxiv.org/abs/1908.08286>. arXiv: 1908.08286.
31. M. Lothaire, editor. *Combinatorics on Words*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2 edition, 1997. ISBN 9780521599245. . URL <https://www.cambridge.org/core/books/combinatorics-on-words/6FEBB4FCCB43895CCEFA8D69A0983374>.
32. Terry J. Lyons, Michael J. Caruana, and Thierry Lévy. *Differential Equations Driven by Rough Paths: Ecole d'Été de Probabilités de Saint-Flour XXXIV-2004*. École d'Été de Probabilités de Saint-Flour. Springer-Verlag, Berlin Heidelberg, 2007. ISBN 978-3-540-71284-8. . URL <https://www.springer.com/gp/book/9783540712848>.
33. Fenglong Ma, Radha Chitta, Jing Zhou, Quanzeng You, Tong Sun, and Jing Gao. Dipole: Diagnosis Prediction in Healthcare via Attention-based Bidirectional Recurrent Neural Networks. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*, pages 1903–1911, 2017. . URL <http://arxiv.org/abs/1706.05764>. arXiv: 1706.05764.
34. James Morrill, Andrey Kormilitzin, Alejo Nevado-Holgado, Sumanth Swaminathan, Sam Howison, and Terry Lyons. The Signature-Based Model for Early Detection of Sepsis from Electronic Health Records in the Intensive Care Unit. In *2019 Computing in Cardiology Conference*, December 2019. . URL <http://www.cinc.org/archives/2019/pdf/CinC2019-014.pdf>.
35. James Morrill, Adeline Fermanian, Patrick Kidger, and Terry Lyons. A Generalised Signature Method for Time Series. *arXiv:2006.00873 [cs, stat]*, June 2020. URL <http://arxiv.org/abs/2006.00873>. arXiv: 2006.00873.
36. Max Pfeffer, Anna Seigal, and Bernd Sturmfels. Learning Paths from Signature Tensors. *arXiv:1809.01588 [cs, math, stat]*, November 2018. URL <http://arxiv.org/abs/1809.01588>. arXiv: 1809.01588.
37. Trang Pham, Truyen Tran, Dinh Phung, and Svetha Venkatesh. DeepCare: A Deep Dynamic Memory Model for Predictive Medicine. In James Bailey, Latifur Khan, Takashi Washio, Gill Dobbie, Joshua Zhexue Huang, and Ruili Wang, editors, *Advances in Knowledge Discovery and Data Mining*, Lecture Notes in Computer Science, pages 30–41, Cham, 2016. Springer International Publishing. ISBN 978-3-319-31750-2. .
38. Zhaozhi Qian, Ahmed M. Alaa, Alexis Bellot, Jem Rashbass, and Mihaela van der Schaar. Learning Dynamic and Personalized Comorbidity Networks from Event Data using Deep Diffusion Processes. *arXiv:2001.02585 [cs, stat]*, January 2020. URL <http://arxiv.org/abs/2001.02585>. arXiv: 2001.02585.
39. Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M. Dai, Nissan Hajaj, Michaela Hardt, Peter J. Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, Patrik Sundberg, Hector Yee, Kun Zhang, Yi Zhang, Gerardo Flores, Gavin E. Duggan, Jamie Irvine, Quoc Le, Kurt Litsch, Alexander Mossin, Justin Tansuwan, De Wang, James Wexler, Jimbo Wilson, Dana Ludwig, Samuel L. Volchenboum, Katherine Chou, Michael Pearson, Srinivasan Madabushi, Nigam H. Shah, Atul J. Butte, Michael D. Howell, Claire Cui, Greg S. Corrado, and Jeffrey Dean. Scalable and accurate deep learning with electronic health records. *npj Digital Medicine*, 1(1):1–10, May 2018. ISSN 2398-6352. . URL <https://www.nature.com/articles/s41746-018-0029-1>.
40. Laila Rasmy, Yonghui Wu, Ningtao Wang, Xin Geng, W. Jim Zheng, Fei Wang, Hulin Wu, Hua Xu, and Degui Zhi. A study of generalizability of recurrent neural network-based predictive models for heart failure onset risk using a large and heterogeneous EHR data set. *Journal of Biomedical Informatics*, 84:11–16, 2018. ISSN 1532-0480. .
41. Jeremy Francis Reizenstein. *Iterated-Integral Signatures in Machine Learning*. PhD thesis, University of Warwick, 2019.
42. Takaya Saito and Marc Rehmsmeier. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLOS ONE*, 10(3):e0118432, March 2015. ISSN 1932-6203. . URL <https://dx.plos.org/10.1371/journal.pone.0118432>.
43. Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, Bette Liu, Paul Matthews, Giok Ong, Jill Pell, Alan Silman, Alan Young, Tim Sprosen, Tim Peakman, and Rory Collins. UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. *PLOS Medicine*, 12(3):e1001779, March 2015. ISSN 1549-1676. . URL <https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1001779>. Publisher: Public Library of Science.
44. Colin Walsh and George Hripcsak. The Effects of Data Sources, Cohort Selection, and Outcome Definition on a Predictive Model of Risk of Thirty-Day Hospital Readmissions. *Journal of biomedical informatics*, 52:418–426, December 2014. ISSN 1532-0464. . URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4261028/>.
45. Shuai Xiao, Junchi Yan, Mehrdad Farajtabar, Le Song, Xiaokang Yang, and Hongyuan Zha. Learning Time Series Associated Event Sequences With Recurrent Point Process Networks. *IEEE Transactions on Neural*

- Networks and Learning Systems*, 30(10):3124–3136, October 2019. ISSN 2162-2388. . Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
46. Zecheng Xie, Zenghui Sun, Lianwen Jin, Hao Ni, and Terry Lyons. Learning Spatial-Semantic Context with Fully Convolutional Recurrent Network for Online Handwritten Chinese Text Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(8):1903–1917, August 2018. ISSN 1939-3539. . Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
 47. Weixin Yang, Terry Lyons, Hao Ni, Cordelia Schmid, and Lianwen Jin. Developing the Path Signature Methodology and its Application to Landmark-based Human Action Recognition. *arXiv:1707.03993 [cs]*, December 2019. URL <http://arxiv.org/abs/1707.03993>. arXiv: 1707.03993.
 48. Jinghe Zhang, Kamran Kowsari, James H. Harrison, Jennifer M. Lobo, and Laura E. Barnes. Patient2Vec: A Personalized Interpretable Deep Representation of the Longitudinal Electronic Health Record. *IEEE Access*, 6: 65333–65346, 2018. ISSN 2169-3536. . URL <http://arxiv.org/abs/1810.04793>. arXiv: 1810.04793.
 49. Jing Zhao, Panagiotis Papapetrou, Lars Asker, and Henrik Boström. Learning from heterogeneous temporal data in electronic health records. *Journal of Biomedical Informatics*, 65:105–119, January 2017. ISSN 1532-0464. . URL <http://www.sciencedirect.com/science/article/pii/S1532046416301654>.