



This is a repository copy of *An efficient representation of chronological events in medical texts*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/168039/>

Version: Published Version

Proceedings Paper:

Kormilitzin, A., Vaci, N. orcid.org/0000-0002-8094-0902, Liu, Q. et al. (3 more authors) (2020) An efficient representation of chronological events in medical texts. In: Proceedings of the 11th International Workshop on Health Text Mining and Information Analysis. OCHI 2020: The 11th International Workshop on Health Text Mining and Information Analysis, 20 Nov 2020, Online. Association for Computational Linguistics , pp. 97-103. ISBN 9781952148811

10.18653/v1/P17

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:
<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

An efficient representation of chronological events in medical texts

Andrey Kormilitzin^{1*}, Nemanja Vaci^{2*}, Qiang Liu¹, Hao Ni^{3,5},
Goran Nenadic^{4,5} and Alejo Nevado-Holgado^{1,6}

¹Department of Psychiatry, University of Oxford, Oxford, OX3 7JX, UK

²Department of Psychology, University of Sheffield, Sheffield, S1 1HD, UK

³Department of Mathematics, University College London, London, WC1H 0AY, UK

⁴School of Computer Science, University of Manchester, Manchester, M13 9PL, UK

⁵The Alan Turing Institute London, London, NW1 2DB, UK

⁶Akrivia Health, Oxford, OX1 1BY, UK

andrey.kormilitzin@psych.ox.ac.uk

Abstract

In this work we addressed the problem of capturing sequential information contained in longitudinal electronic health records (EHRs). Clinical notes, which is a particular type of EHR data, are a rich source of information and practitioners often develop clever solutions how to maximise the sequential information contained in free-texts. We proposed a systematic methodology for learning from chronological events available in clinical notes. The proposed methodological *path signature* framework creates a non-parametric hierarchical representation of sequential events of any type and can be used as features for downstream statistical learning tasks. The methodology was developed and externally validated using the largest in the UK secondary care mental health EHR data on a specific task of predicting survival risk of patients diagnosed with Alzheimer's disease. The signature-based model was compared to a common survival random forest model. Our results showed a 15.4% increase of risk prediction AUC at the time point of 20 months after the first admission to a specialist memory clinic and the signature method outperformed the baseline mixed-effects model by 13.2 %.

1 Introduction

Electronic health records (EHRs) have now become ubiquitous and offer novel opportunities for clinical research by supporting the development of intelligent decision support systems and improvement of patients' care. One of the distinct features of EHR is that the data are being collected over time and might be seen as health data streams, allowing research to study longitudinal trends and make inference about the progression of disease, treatments and outcomes. However, the proper representation of sequential medical events still remains

a challenge. Moreover, longitudinal clinical notes exhibit a multi-level hierarchical structure, where events are described and embedded in sentences, sentences in paragraphs and eventually resulting in chronologically ordered documents. Recent works have addressed the problem of capturing this information directly from raw texts by introducing novel neural network architectures, such as attention-based recurrent neural networks (Bai et al., 2018) and time-aware Transformers (Zhang et al., 2020). When dealing with chronological clinical notes, practitioners make multiple decisions on how to structure and transform these sequential events, which are often simplifications of medical histories. In this work we proposed a different methodology to address the problem of learning from events found in clinical notes, by first extracting them using natural language processing and then representing the sequential order by means of the *path signatures*. The signature (Lyons, 2014) is a non-parametric representation of heterogeneous sequential data, offers a feature extraction method from longitudinal events and can naturally be integrated within a general data mining pipeline. To demonstrate the methodology, we used the largest secondary care mental health EHR data in the UK to develop a survival prognostic model for patients diagnosed with Alzheimer's disease.

2 Method

2.1 Data

The data in this study were sourced from the UK-Clinical Record Interactive Search system (UK-CRIS), which provides a research platform (<https://crisnetwork.co/>) for data mining and analysis using de-identified real-world observational electronic patients records from twelve secondary care UK Mental Health NHS Trusts (Goodday et al., 2020). UK-CRIS provides access to struc-

¹Equal contribution.

tured information, such as ICD-10 coded diagnoses, quality of life scales and demographic information, as well as various unstructured texts, such as clinical summaries, discharge letters and progress notes. The study cohort jointly comprised records from 24,108 patients diagnosed with Alzheimer’s disease and various types of dementia, containing more than 3.7 million individual clinical documents from two centres: Oxford and Southern Health Foundation NHS Trusts. The field of clinical NLP in general, and of mental health and Alzheimer’s research in particular, largely suffers from the dearth of gold-annotated data. The reason is due to the shortage of trained annotators with clinical background who are also authorised to access sensitive patient-level data. Therefore, to develop a robust information extraction (IE) model from an insufficient amount of data, we leveraged the idea of transfer learning using the publicly available MIMIC-III corpus (Johnson et al., 2016) comprising information relating to patients admitted to intensive care units (ICU) with more than 2.1 million clinical notes as well as 505 gold-annotated by clinical experts discharge summaries from the 2018 n2c2 challenge (Henry et al., 2020). We assert that the study was independently approved and granted by the Oxfordshire and Southern Health NHS Foundation Trust Research Ethics Committees.

2.2 Information extraction model

The information extraction model was developed to identify diagnosis, medications and cognitive health assessment Mini-Mental State Examination score (MMSE) (Pangman et al., 2000). Additionally, the identified entities were classified according to several attributes, such as the ‘experiencer’ modality (i.e., whether the MMSE was actually referring to a patient or to a family member), temporal information (i.e the date of diagnosis or MMSE score) and negations (i.e. discontinued medications) (Harkema et al., 2009). Such drug mentions were discarded in order to extract the most accurate information. Generic and brand drug names were normalised using the British National Formulary, the core pharmaceutical reference book (Committee et al., 2019). The architecture of the named entity recognition model comprised a hybrid approach of an ontology-based fuzzy pattern matching and a bi-directional LSTM neural network architecture with the attention mechanism (Bahdanau et al., 2014) for sequence classification.

The GloVe word embedding (Pennington et al., 2014) were fine-tuned on both MIMIC-III and UK-CRIS data (Vaci et al., 2020; Kormilitzin et al., 2020). The developed IE model was trained only on data from the Oxford Health NHS Trust instance and externally validated on a sample of data from a regionally different Southern Health NHS Foundation Trust.

2.3 The signature of a path

Repeated measurements, speech, text, time-series or any other sequential data might be seen as a path-valued random variable. Formally, a path X of finite length in d dimensions can be described by the mapping $X : [a, b] \rightarrow \mathbb{R}^d$, or in terms of co-ordinates $X = (X_t^1, X_t^2, \dots, X_t^d)$, where each coordinate X_t^i is real-valued and parametrised by $t \in [a, b]$. The signature representation S of a path X is defined as an infinite series:

$$S(X)_{a,b} = (1, S(X)_{a,b}^1, S(X)_{a,b}^2, \dots, S(X)_{a,b}^d, S(X)_{a,b}^{1,1}, S(X)_{a,b}^{1,2}, \dots), \quad (1)$$

where each term is a k -fold iterated integral of the path X labelled by multi-index i_1, \dots, i_k :

$$S(X)_{a,b}^{i_1, \dots, i_k} = \int_{a < t_k < b} \dots \int_{a < t_1 < t_2} dX_{t_1}^{i_1} \dots dX_{t_k}^{i_k}. \quad (2)$$

However, in many real-life applications the first k -terms of the truncated signature at level L give a satisfying approximation. Intuitively, it is analogous to statistical moments of a d -dimensional vector-valued random variable, such as mean, variance or higher moments. One can define statistical moments of a *path*-valued random variable, which are essentially the *signature* moments (Chevyrev and Oberhauser, 2018) defined in Eq. (2). The signature $S(X)$ completely characterises a path X up to tree-like equivalence and is invariant to reparameterisation (Hambly and Lyons, 2010). The signature can also be expressed in a more compact form known as *log-signature* (Liao et al., 2019; Morrill et al., 2020a), which is the formal power series of $\log S(X)$, while carrying the same information. Informally, the path signature captures the order of events. For example, consider two sequences $X_1 = aabba$ and $X_2 = baaab$ consisting of a simple vocabulary with only two letters $\{a, b\}$. The sequences might be presented as paths in 2d space as shown in Fig. 1. Each linear segment between




Figure 1: Two paths $X_1 = aabba$ and $X_2 = baaab$.

two points (Fig. 1) corresponds to a single letter in the sequence and the arrows denote the temporal direction of the sequence. Despite the same

Level	1	2	3	4
$S(X_1)$	3 2	1	-0.5 -1	-1/3 -0.5 0
$S(X_2)$	3 2	0	1.5 0.5	0 0 0

Table 1: The first $k = 8$ terms of the log signature expansion up to level $L = 4$. The difference between two sequences X_1 and X_2 is apparent starting from the second level.

number of letters in the sequences $\{a = 3, b = 2\}$, the order of letters matters. The signature easily picks the differences and the first four levels of the log-signatures of paths are shown in Table 1. The lower order signature terms $S^{(i)}$ are the increments along the i -th direction (i.e. the distance between the endpoints), for example, $S^{(1)} = 3 - 0 = 3$ and $S^{(2)} = 2 - 0 = 2$ as can be seen in Figure 1. The second order corresponds to the area enclosed by a path and a chord connecting endpoints (Chevyrev and Kormilitzin, 2016).

The usefulness of a path signature as a feature map of sequential data was demonstrated theoretically (Chevyrev and Oberhauser, 2018) as well as in numerous machine learning applications in healthcare (Morrill et al., 2019; Kormilitzin et al., 2016; Arribas et al., 2018; Morrill et al., 2020b; Kormilitzin et al., 2017), finance (Arribas, 2018), computer vision (Yang et al., 2017; Xie et al., 2017), topological data analysis (Chevyrev et al., 2018) and deep learning (Kidger et al., 2019).

2.4 Independent and outcome variables

The independent variables used in the prognostic model were medications and the MMSE scores collected over time. The dependent outcome variable was right-censored time to death data in months. A synthetic example of the patient’s records (Table 2) and the corresponding algorithmically extracted longitudinal data is presented in Table 3.

The outcome variable was encoded as a tuple: $(True, 34.17)$ indicating that a person has died after 34.17 months since the very first visit to a specialist memory clinic. The patient was treated by two different medications with a changing pattern and eventually was tapered off medication due to no further expected improvement.

2.5 Baseline longitudinal data summarisation

The signature transformation might be seen as a hierarchical statistical summarisation (“feature extraction”) of the longitudinal data along the temporal dimension. In order to benchmark the proposed method, we used a time-honoured linear mixed-effects regression as a baseline model for longitudinal summarisation. Specifically, each patient-level longitudinal MMSE scores were modelled using a linear regression and the resulting coefficients, such as an intercept and a slope, were used as features representing the progression of the MMSE over time. The median number of medication categories was used as an additional feature, resulting in three features for each patient.

2.6 Survival random forests

The common statistical approach to analyse the time-to-event survival data is based on the linear Cox model (Collett, 2015). However, Miao et al. (2015) showed that a survival random forest (SRF) approach (Ishwaran et al., 2008) outperformed linear Cox model, based on the Harrell’s concordance index (C-index) (Harrell et al., 1982), and was understandably capable of identifying non-linear effects of the input variables as opposed to linear Cox model. Therefore, we chose the SRF as the preferred method. The SRF approach was implemented in Python using “scikit-survival” package (Pölsterl et al., 2015). The Harrell’s C-index (the concordance index) is a goodness of fit measure for risk scores models. It is a common statistical approach to evaluate risk models in survival analysis, where data may be right-censored and corresponds to rank correlation between predicted risk scores and observed time points, similarly to Kendall’s τ .

3 Results

3.1 Information extraction model

We used a hybrid approach to developing an IE model consisting of training a baseline model using MIMIC-III and n2c2 annotated data. Specifically, the named-entity recognition (NER) model

Doc date	Text
05-Oct-2016	<i>Today I saw a patient diagnosed with Alzheimer's, who deteriorated: MMSE 23/30 as compared to 25/30 from 1st January. Started on Rivastigmine.</i>
12-Feb-2017	<i>Today MMSE 19, the patient didn't respond to Rivastigmine and was changed to Donepezil.</i>
03-Feb-2018	<i>Great response to new treatment (MMSE 23/30), continue on Donepezil.</i>
01-Apr-2019	<i>The patient stopped responding to Donepezil and severely deteriorated (MMSE 14/30), stop Donepezil.</i>

Table 2: A synthetic example of chronological medical records.

Date	Medication	MMSE
01-Jan-2016	NoMed	25/30
05-Oct-2016	Rivastigmine	22/30
12-Feb-2017	Donepezil	19/30
03-Feb-2018	Donepezil	23/30
01-Apr-2019	Discontinued	14/30

Table 3: Extracted and chronologically structured data from Table 2.

comprised a transition-based system based on the chunking model (Lample et al., 2016) where tokens were represented as hashed and embedded representations of the prefix, suffix, shape and lemmatised features of words, followed by the rule-based matching using the BNF vocabulary. The IE model was implemented using “spaCy” python library¹, including negations and temporal information identification as well as relationships classification between the word-tokens using linguistic features, such as part-of-speech and dependencies. Finally, the active learning tool “Prodigy”² was used for iterative model improvement. Target domain training, validation and external validation data contained a collection of gold-annotated drug names, diagnosis and cognitive health assessment MMSE scores as shown in Table 4. The IE model

Concept	Training	Validation	External val.	Total
Drug	216	153	30	399
MMSE	169	87	23	279
Diagnosis	570	352	26	948

Table 4: The number of gold-annotated instances in the training, validation and external validation data sets.

achieved a good and consistent performance on both validation and external validation data sets (Table 5). The annotation schema was developed following the recommendations of Pustejovsky and Stubbs (2012). The token-level performance metrics were evaluated using the SemEval schema (Segura Bedmar et al., 2013) and the inter-annotator agreement (IAA) of two clinical annotators was

¹<https://spacy.io>

²<https://prodi.gy>

computed using F1 score.

Concept	Validation			External val.			IAA
	Pr	Re	F1	Pr	Re	F1	F1
Diagnosis	89.6	96.3	92.8	84.1	86.3	84.8	95
Drug	98	98	98.1	92.4	68.4	78.3	96
MMSE	92.6	74.7	82.8	85.6	81.2	82.6	100

Table 5: Performance (shown in %) of the information extraction model. IAA - inter annotator agreement.

	n	male	female	survival time
died	1962	841	1121	52.2(22.8)
censored	1500	529	971	28.4(16.6)

Table 6: Summary statistics of the extracted data for survival analysis. Survival time is shown as mean(std) in months. The MMSE scores were not observed for censored people later in time, while date of death was recorded in hospital.

3.2 Prognostic model

Four prognostic models were developed and compared to each other. All models estimated the survival probability of a patient diagnosed with Alzheimer’s disease since their first admission to a memory clinic. We compared signature (“Sig”, Sec. 2.3) versus non-signature (“Non-sig”, Sec. 2.5) models. We also estimated the added value of the sequential information contained in the treatment course with medications. Specifically, we used two sets of input variables: {time, MMSE} and {time, MMSE, medications}, where time corresponds to the date of MMSE score or prescribed medication as presented in Table 3. For the “Sig” model, the input variable were first transformed into signatures, where the categorical medication names were one-hot encoded and augmented with numerical MMSE scores to create a path. For the “Non-sig” model, the longitudinal MMSE scores were summarised by means of linear models adjusting for each patients and the median number of distinct medications were computed. Both models were trained and validated using the same folds of stratified 5-fold cross validation (with fixed random seed). The quality of predictions was assessed

using the Harell’s C-index and the results are summarised in Table 7. The signatures were computed using the “esig” Python library³, however, alternative libraries are also available (Reizenstein and Graham, 2018; Kidger and Lyons, 2020).

Features	Sig	Non-sig
{time, MMSE}	0.626(0.009)	0.574(0.022)
{time, MMSE, meds}	0.621(0.011)	0.571(0.019)

Table 7: Harell’s C-index measure of four models. Values reported as mean(std) over 5-fold cross validation.

We also estimated the time-dependent area under the curve of receiver operating characteristics (Lambert and Chevret, 2016). It is a natural extension of a common AUC ROC analysis to possibly censored survival times where the patients’ cognitive health is usually better at the very first visit to a memory clinic, while their condition may deteriorate later. The time-dependent cumulative dynamic AUC ROC of all four models are presented in Fig. 2. The signature features outperformed the non-signature ones at all times and the inclusion of sequential information from switching medications improved AUC ROC at later times. However, both models struggle to reliably predict the future outcomes further than 3 years. This is due to the limitation of predictors and the available number of patients after 3 years rather than the capacity of our model.

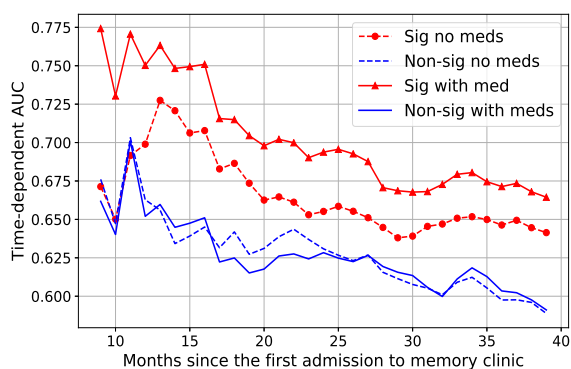


Figure 2: Time-dependent AUC of risk prediction over time since the first admission to a memory clinic.

4 Discussion and future direction

Unstructured longitudinal electronic health records, such as free-text clinical notes, inherently contain rich information about patients’ health and outcomes over time. The right analytical tools capable

³<https://esig.readthedocs.io/>

of capturing sequential information can therefore maximise utilisation of longitudinal EHRs and can be valuable for supporting clinical decisions and prognostic models. In this work we implemented a signature-based approach to represent chronological events extracted using natural language processing from clinical notes. Extracted chronological events can be seen as a trajectory (path) embedded in a high-dimensional multi-modal space of events (i.e. different medications, interventions, measures, etc) and the signature uniquely characterises the path in the most succinct way. The signature-based feature extraction approach was compared to hand-crafted features, comprising a slope and an intercept of MMSE scores over time and the median number of medications for each patient. The signatures represent a hierarchical collection of features, where the first order is proportional to linear statistical moments (i.e. mean) and is not sensitive to the order of data points, as illustrated in Table 1. We demonstrated that the sequential information about medications has significantly improved the time-dependent AUC as captured by the signatures (Figure 2). In future works we will extend the proposed framework to include the structured information available in EHR (i.e. lab results, coded procedures or clinical encounters) and will develop an interpretability framework to make the signature-based models explainable.

Acknowledgments

The study was funded by the National Institute for Health Research’s (NIHR) Oxford Health Biomedical Research Centre (BRC-1215-20005). This work was supported by the UK Clinical Records Interactive Search (UK-CRIS) system funded and developed by the NIHR Oxford Health BRC at Oxford Health NHS Foundation Trust and the Department of Psychiatry, University of Oxford. AK, NV, QL, ANH are funded by the MRC Pathfinder Grant (MC-PC-17215). HN was supported by the Alan Turing Institute under the EP-SRC grant EP/N510129/1 and under the EPSRC EP/S026347/1. The views expressed are those of the authors and not necessarily those of the UK National Health Service, the NIHR, or the UK Department of Health. We highly appreciate the work and support of the Oxford CRIS Team: Tanya Smith, Lulu Kane, Adam Pill and Suzanne Fisher and Dr Peter Phiri of the Southern Health NHS Foundation Trust CRIS team.

References

- Imanol Perez Arribas. 2018. Derivatives pricing using signature payoffs. *arXiv preprint arXiv:1809.09466*.
- Imanol Perez Arribas, Kate Saunders, Guy Goodwin, and Terry Lyons. 2018. A signature-based machine learning model for bipolar disorder and borderline personality disorder. *Translational Psychiatry*, 8.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Tian Bai, Shanshan Zhang, Brian L Egleston, and Slobodan Vucetic. 2018. Interpretable representation learning for healthcare via capturing disease progression through time. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 43–51.
- Ilya Chevyrev and Andrey Kormilitzin. 2016. A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*.
- Ilya Chevyrev, Vidit Nanda, and Harald Oberhauser. 2018. Persistence paths and signature features in topological data analysis. *IEEE transactions on pattern analysis and machine intelligence*, pages 1–1.
- Ilya Chevyrev and Harald Oberhauser. 2018. Signature moments to characterize laws of stochastic processes. *arXiv preprint arXiv:1810.10971*.
- David Collett. 2015. *Modelling survival data in medical research*. CRC press.
- Joint Formulary Committee et al. 2019. *BNF 77 (British National Formulary) March 2019*. Pharmaceutical Press.
- SM Goodday, A Kormilitzin, N Vaci, Q Liu, A Cipriani, T Smith, and A Nevado-Holgado. 2020. Maximizing the use of social and behavioural information from secondary care mental health electronic health records. *Journal of Biomedical Informatics*, page 103429.
- Ben Hambly and Terry Lyons. 2010. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, pages 109–167.
- Henk Harkema, John N Dowling, Tyler Thornblade, and Wendy W Chapman. 2009. Context: an algorithm for determining negation, experiencer, and temporal status from clinical reports. *Journal of biomedical informatics*, 42(5):839–851.
- Frank E Harrell, Robert M Califf, David B Pryor, Kerry L Lee, and Robert A Rosati. 1982. Evaluating the yield of medical tests. *Jama*, 247(18):2543–2546.
- Sam Henry, Kevin Buchan, Michele Filannino, Amber Stubbs, and Ozlem Uzuner. 2020. 2018 n2c2 shared task on adverse drug events and medication extraction in electronic health records. *Journal of the American Medical Informatics Association*, 27(1):3–12.
- Hemant Ishwaran, Udaya B Kogalur, Eugene H Blackstone, Michael S Lauer, et al. 2008. Random survival forests. *The annals of applied statistics*, 2(3):841–860.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3:160035.
- Patrick Kidger, Patric Bonnier, Imanol Perez Arribas, Christopher Salvi, and Terry Lyons. 2019. Deep signature transforms. In *Advances in Neural Information Processing Systems*, pages 3105–3115.
- Patrick Kidger and Terry Lyons. 2020. Signatory: differentiable computations of the signature and logsignature transforms, on both cpu and gpu. *arXiv preprint arXiv:2001.00706*.
- AB Kormilitzin, KEA Saunders, PJ Harrison, JR Geddes, and TJ Lyons. 2016. Application of the signature method to pattern recognition in the cequel clinical trial. *arXiv preprint arXiv:1606.02074*.
- Andrey Kormilitzin, Kate EA Saunders, Paul J Harrison, John R Geddes, and Terry Lyons. 2017. Detecting early signs of depressive and manic episodes in patients with bipolar disorder using the signature-based model. *arXiv preprint arXiv:1708.01206*.
- Andrey Kormilitzin, Nemanja Vaci, Qiang Liu, and Alejo Nevado-Holgado. 2020. Med7: a transferable clinical natural language processing model for electronic health records. *arXiv preprint arXiv:2003.01271*.
- Jérôme Lambert and Sylvie Chevret. 2016. Summary measure of discrimination in survival models based on cumulative/dynamic time-dependent roc curves. *Statistical methods in medical research*, 25(5):2088–2102.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. *arXiv preprint arXiv:1603.01360*.
- Shujian Liao, Terry Lyons, Weixin Yang, and Hao Ni. 2019. Learning stochastic differential equations using rnn with log signature features. *arXiv preprint arXiv:1908.08286*.
- Terry Lyons. 2014. Rough paths, signatures and the modelling of functions on streams. *arXiv preprint arXiv:1405.4537*.

- Fen Miao, Yun-Peng Cai, Yuan-Ting Zhang, and Chun-Yue Li. 2015. Is random survival forest an alternative to cox proportional model on predicting cardiovascular disease? In *6TH European conference of the international federation for medical and biological engineering*, pages 740–743. Springer.
- James Morrill, Patrick Kidger, Cristopher Salvi, James Foster, and Terry Lyons. 2020a. Neural cdes for long time series via the log-ode method. *arXiv preprint arXiv:2009.08295*.
- James Morrill, Andrey Kormilitzin, Alejo Nevado-Holgado, Sumanth Swaminathan, Sam Howison, and Terry Lyons. 2019. The signature-based model for early detection of sepsis from electronic health records in the intensive care unit. In *2019 Computing in Cardiology (CinC)*, pages Page–1. IEEE.
- James H Morrill, Andrey Kormilitzin, Alejo J Nevado-Holgado, Sumanth Swaminathan, Samuel D Howison, and Terry J Lyons. 2020b. Utilization of the signature method to identify the early onset of sepsis from multivariate physiological time series in critical care monitoring. *Critical Care Medicine*, 48(10):e976–e981.
- Verna C Pangman, Jeff Sloan, and Lorna Guse. 2000. An examination of psychometric properties of the mini-mental state examination and the standardized mini-mental state examination: implications for clinical practice. *Applied Nursing Research*, 13(4):209–213.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Sebastian Pölsterl, Nassir Navab, and Amin Katouzian. 2015. Fast training of support vector machines for survival analysis. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 243–259. Springer.
- James Pustejovsky and Amber Stubbs. 2012. *Natural Language Annotation for Machine Learning: A guide to corpus-building for applications*. ” O’Reilly Media, Inc.”.
- Jeremy Reizenstein and Benjamin Graham. 2018. The iisignature library: efficient calculation of iterated-integral signatures and log signatures. *arXiv preprint arXiv:1802.08252*.
- Isabel Segura Bedmar, Paloma Martínez, and María Herrero Zazo. 2013. Semeval-2013 task 9: Extraction of drug-drug interactions from biomedical texts (ddiextraction 2013). Association for Computational Linguistics.
- Nemanja Vaci, Ivan Koychev, Chi-Hun Kim, Andrey Kormilitzin, Qiang Liu, Christopher Lucas, Azad Dehghan, Goran Nenadic, and Alejo Nevado-Holgado. 2020. Real-world effectiveness, its predictors and onset of action of cholinesterase inhibitors and memantine in dementia: retrospective health record study. *The British Journal of Psychiatry*, page 1–7.
- Zecheng Xie, Zenghui Sun, Lianwen Jin, Hao Ni, and Terry Lyons. 2017. 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.
- Weixin Yang, Terry Lyons, Hao Ni, Cordelia Schmid, Lianwen Jin, and Jiawei Chang. 2017. Leveraging the path signature for skeleton-based human action recognition. *arXiv preprint arXiv:1707.03993*.
- Dongyu Zhang, Jidapa Thadajarassiri, Cansu Sen, and Elke Rundensteiner. 2020. Time-aware transformer-based network for clinical notes series prediction. In *Machine Learning for Healthcare Conference*, pages 566–588. PMLR.