

Information Extraction from Swedish Medical Prescriptions with Sig-Transformer Encoder

John Pougué Biyong¹ Bo Wang^{2,3} Terry Lyons^{1,3} Alejo J Nevado-Holgado²

¹Mathematical Institute, University of Oxford, UK

²Department of Psychiatry, University of Oxford, UK

³The Alan Turing Institute, London, UK

john.pougue-biyong@maths.ox.ac.uk, bo.wang@psych.ox.ac.uk

Abstract

Relying on large pretrained language models such as Bidirectional Encoder Representations from Transformers (BERT) for encoding and adding a simple prediction layer has led to impressive performance in many clinical natural language processing (NLP) tasks. In this work, we present a novel extension to the Transformer architecture, by incorporating signature transform with the self-attention model. This architecture is added between embedding and prediction layers. Experiments on a new Swedish prescription data show the proposed architecture to be superior in two of the three information extraction tasks, comparing to baseline models. Finally, we evaluate two different embedding approaches between applying Multilingual BERT and translating the Swedish text to English then encode with a BERT model pretrained on clinical notes.

1 Introduction

Medical prescription notes written by clinicians about patients contains valuable information that the structured part of electronic health records (EHRs) does not have. Manually extracting and annotating such information from large amount of textual fields in medical records such as prescription notes is time-consuming, and expensive since it has to be provided by domain specialists who are in high demand. Information extraction (IE), a specialised area in natural language processing (NLP), refers to the automatic extraction of structured information such as concepts, entities and events from free text (Wang et al., 2018). While rule-base IE systems are still widely used in healthcare (Banda et al., 2018; Davenport and Kalakota, 2019), machine learning-based data-driven approaches have gained much more interest due to their ability to scale and learn to recognise complex patterns. In this paper, we present a set of medical prescriptions

written in Swedish and develop a neural network-based system that extracts important information from the textual data.

Recent advances in pretraining large-scale contextualised language representations, including ELMo (Peters et al., 2018) and BERT (Devlin et al., 2018), have proven to be a successful strategy for transfer learning and pushed the performance in many NLP tasks of general purpose. Clinical notes differ substantially to general-domain text and biomedical literature, in terms of its linguistic characteristics. Several studies have fine-tuned general-domain language models such as BERT on in-domain clinical text (e.g. electronic health records) (Si et al., 2019; Peng et al., 2019; Alsentzer et al., 2019; Huang et al., 2020) for downstream clinical NLP tasks. As summarised in (Gu et al., 2020) many such tasks can be formulated as a classification or regression task, wherein either a simple linear layer or sequential models such as LSTM or conditional random field (CRF) are added after the BERT encoding of the input text as task-specific prediction layer.

BERT’s model architecture is based on the Transformer model (Vaswani et al., 2017), which employs the self-attention mechanism to attend to different positions of the input sequence. As it doesn’t contain any recurrence or convolution, the model has to add positional encoding in order to model the order of the sequence (e.g. word order). Signature transform, initially introduced in rough path theory as a branch of stochastic analysis, is a non-parametric approach of encoding sequential data while capturing the order information in the data. It has shown successes in the recent years in a range of machine learning tasks involving sequential modelling (Arribas et al., 2018; Morrill et al., 2020; Toth and Oberhauser, 2020; Wang et al., 2020). We hypothesise the signature transform method can be integrated in the Transformer

model, in which its ability to naturally capture sequential ordering complements the ability of attending to important positions and learning long-range dependencies in self-attention.

As a preliminary study, our aim is to develop a machine learning-based system that extracts relevant information from the Swedish medical prescription notes¹, namely *quantity*, *quantity tag* and *indication*. The contributions of this work are as follows:

1. We experiment two different approaches of embedding the Swedish prescription notes. One encodes the Swedish text directly using Multilingual BERT (M-BERT) (Devlin et al., 2018) while for the other approach we translate the prescriptions and then apply ClinicalBERT (Huang et al., 2020) that is pretrained on clinical text in English.
2. We propose an extension to the Transformer model, named Sig-Transformer Encoder (STE), which integrates signature transform into the Transformer architecture so the order information in the prescription notes can be learnt in a more effective way. To the best of our knowledge, this is the first attempt to integrate signature transform in the Transformer architecture.
3. We demonstrate good performance in two of the three tasks, namely *quantity* and *quantity tag*. As for the *indication* task, we provide an analysis on why the models fail.
4. We show one of our proposed STE-based approaches, namely M-BERT+STE, outperforming other baseline models.

2 Background

Clinical language representation has attracted an increasing interest in the natural language processing (NLP) community (Kalyan and Sangeetha, 2020). More recently with the advent of Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2018), fine-tuning of general-domain language models has been widely adopted for many clinical NLP tasks². Among them two clinicalBERT studies (Alsentzer et al., 2019; Huang

¹Our ultimate goal is to apply such model to a much larger database. However, this is beyond the scope of this paper.

²Gu et al. (2020) have challenged this strategy and reported better performance for PubMed-based biomedical tasks by conducting domain-specific pretraining from scratch. How-

et al., 2020) have conducted fine-tuning from either BERT or BioBERT (Lee et al., 2020) using de-identified clinical notes from MIMIC-III (Johnson et al., 2016). In particular, Huang et al. (2020) showed the effectiveness of its ClinicalBERT-based model for predicting hospital readmissions.

Many NLP applications in the medical domain can be formulated as either token classification, sequence classification or sequence regression, in which a pretrained language model such as clinicalBERT can be used to encode the input token sequence (which returns either the encoding for every input token or just the [CLS] token) and then a task-specific prediction model is added on top to generate the final output (Gu et al., 2020). In this section, we describe the Transformer encoder and the signature transform, which are the bases of our proposed architecture.

2.1 Transformer Encoder

The Transformer model employs an encoder-decoder structure (Vaswani et al., 2017). Its encoder, which is the model architecture of BERT, is composed of a stack of N identical layers. Each layer has a multi-head self-attention mechanism followed by a position-wise fully connected feed-forward network. The attention function used in the Transformer model takes three input vectors: query (Q), key (K) and value (V). It generates an output vector by computing the weighted sum of the values. The weights are computed by the dot products of the query and all keys, scaled and applied a softmax function.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

Multi-head attention splits Q , K and V into multiple heads by linearly projection, which allows the model to jointly attend to information at different positions from different representation subspaces. Each projected head goes through the scaled dot-product attention function, then concatenated and projected again to output the final values. The Transformer encoder does not explicitly model position information in its structure, and instead it requires adding representations of absolute positions (i.e. positional encoding) to its inputs. Shaw et al. (2018) presented an extension to self-attention, which incorporates relative position information

ever, it did not address the clinical domain and has left it for future work.

for the input data. In this work, we propose to integrate signature transform with self-attention given that path signatures have been proven to be an effective way of capturing the sequential order information in the data.

2.2 Signature Transform

The theory of rough paths, developed by Lyons (1998), can be thought of as a non-linear extension of the classical theory of controlled differential equations. *Path signature* or simply *signature*, is an infinite collection of statistics characterising the underlying path (a discretised version of a continuous path), and *signature transform* is the map from a path to its signature³. Consider a d -dimensional time-dependent path P over the time interval $[0, T] \subset \mathbb{R}$, to a continuous map $P : [0, T] \rightarrow \mathbb{R}^d$. The signature $S(P)$ of this path P over time interval $[0, T]$ is the infinite collection of all iterated integrals of P such that every continuous function of the path may be approximated arbitrarily well by a linear function of its signature:

$$S(P)_{0,T} = (1, S(P)_{0,T}^1, \dots, S(P)_{0,T}^d, S(P)_{0,T}^{1,1}, \dots)$$

where the 0^{th} term is 1 by convention, and the superscripts of the terms after the 0^{th} term run along the set of all multi-index $\{(i_1, \dots, i_k) | k \geq 1, i_1, \dots, i_k \in [d]\}$ with the coordinate iterated integral being:

$$S(P)_{0,T}^{i_1, \dots, i_k} = \int_{t_1, \dots, t_k \in [0, T]} \dots \int_{t_1 < \dots < t_k} dP_{t_1}^{i_1} \otimes \dots \otimes dP_{t_k}^{i_k}$$

where $\forall k \geq 1, P_t \in \mathbb{R}^d, \forall t \in [0, T]$. $S(P)_{0,T}^{i_1, \dots, i_k}$ is termed as the k th level of the signature. In practice we truncate the signature to order n , where the degree of its iterated integrals is no greater than n . This ensures the path signature has finite dimensional representation. Let $TS(P)_{0,T}^n$ denote the truncated signature of P of order n , i.e.

$$TS(P)_{0,T}^n = (1, S(P)_{0,T}^1, \dots, S(P)_{0,T}^{k_n})$$

Therefore the dimensionality of the truncated path signature is $(d^{n+1} - d)(d - 1)^{-1}$. We describe signature transform in more (mathematical) details in Appendix A.2.

³We refer the reader to (Lyons, 2014) for a rigorous introduction of signature transform, and (Chevyrev and Kormilitzin, 2016) for a primer on its use in machine learning.

In recent years using path signatures as features in a suitable neural network model has shown success in various applications, such as online handwritten Chinese character recognition (Yang et al., 2016; Xie et al., 2018), action recognition in videos (Yang et al., 2017) and speech emotion recognition (Wang et al., 2019). More recently, Kidger et al. (2019) proposed to use signature transform deeper within a network, rather than as a feature transformation. Kidger and Lyons (2020) developed differentiable computation of signature transform on both CPU and GPU. To the best of our knowledge, we are the first to incorporate signature transform in the Transformer architecture.

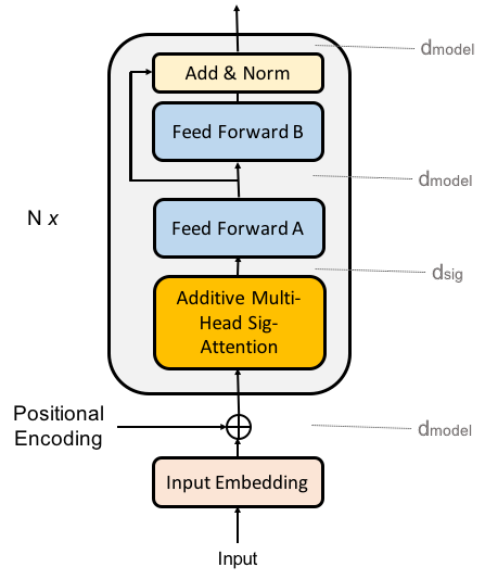


Figure 1: The Sig-Transformer Encoder. d_{sig} is the dimension size in the truncated signature, and d_{model} is the dimension size in the embedding layer as well as before and after the second feed-forward sub-layer.

3 Sig-Transformer Encoder

In this section, we present our proposed extension to the Transformer Encoder (TE), named Sig-Transformer Encoder (STE). As depicted in Figure 1, STE consists of a stack of N identical layers, and each layer has three sub-layers. It replaces the multi-head self-attention sub-layer in TE with our Additive Multi-Head Sig-Attention mechanism, followed by two position-wise fully connected feed-forward sub-layers. Same as the Transformer model, we also employ a residual connection (He et al., 2016) at the last sub-layer, followed by layer normalisation (Ba et al., 2016). The two feed-forward sub-layers and the embed-

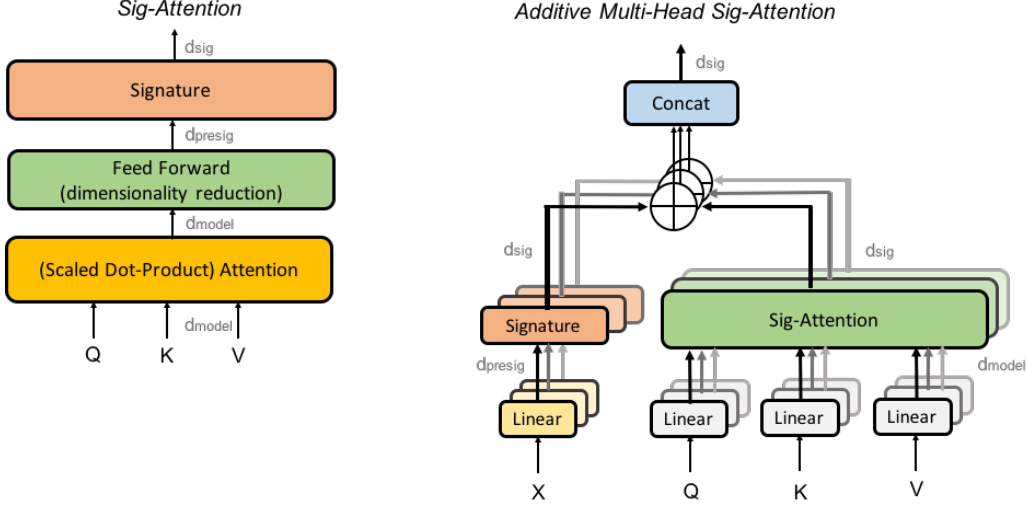


Figure 2: (left) Sig-Attention mechanism; (right) Additive Multi-Head Sig-Attention.

ding layers, produce output of the same dimension $d_{model} = 768$.

3.1 Sig-Attention

The core idea behind our proposed Additive Multi-Head Sig-Attention is the Sig-Attention mechanism (Figure 2 left). It takes the query (Q), key (K) and value (V) vectors as input, and performs three consecutive transformations: (1), scaled dot-product attention; (2), position-wise fully linear map, for dimensionality reduction; (3), signature transform. Its output is computed as follows:

$$\text{SigAttention}(Q, K, V) = \text{ReducedSig}(\text{Attention}(Q, K, V))$$

where

- $\text{Attention}(Q, K, V)$ is as defined in Equation (1);
- $\text{ReducedSig}(x) = S^N(xW + b)$ for any x ;
- S^N is the signature truncated at a given order N ;
- W and b are respectively the weight and bias matrices of the linear map.

As the size of the truncated signature increases exponentially with its input dimension, it is important to reduce the size of attended context vector before applying signature transform. We refer the readers to Table 5 for examples of how the size of the truncated signature (d_{sig}) is determined by its input dimension (d_{presig}) and the order of signature truncation ($order_{sig}$).

3.2 Additive Multi-Head Sig-Attention

In the proposed Additive Multi-Head Sig-Attention model, we combine the information from the input sequence with the output of Sig-Attention, in different representation subspaces (Figure 2 right). The model takes the embedding of an input sequence X as well as the queries, keys and values as input:

$$\begin{aligned} \text{AdditiveMultiHead}(X, Q, K, V) = \\ \text{Concat}(\text{head}_1, \dots, \text{head}_h) \\ \text{where } \text{head}_i = \text{ReducedSig}(X) + \\ \text{SigAttention}(QW_i^Q, KW_i^K, VW_i^V) \end{aligned}$$

and h is the number of parallel heads, the projections are parameter matrices $W_i^Q \in \mathbb{R}^{d_{model} \times d_k}$, $W_i^K \in \mathbb{R}^{d_{model} \times d_k}$, $W_i^V \in \mathbb{R}^{d_{model} \times d_k}$ and $d_k = d_{model}/h$.

As shown in Figure 2 (right), we linearly project the queries, keys and values as well as the input embeddings to d_{model} and d_{presig} dimensions. Then the query, key and value vectors will be fed into the Sig-Attention mechanism, while having been reduced its dimensions by the linear projection the input vector X will be applied to signature transform. The two output signatures are added, resulting in one signature vector $s_i \in \mathbb{R}^{d_{sig}}$ for each head i . Finally all the signature vectors are concatenated and passed onto the next sub-layer.

3.3 Position-wise Feed-Forward Networks

The second sub-layer in Figure 1 is a linear transformation $W \in \mathbb{R}^{d_{sig} \times d_{model}}$ which is applied to each position separately and identically. It simply brings

back the dimensions from d_{sig} to d_{model} . The third sub-layer is a ReLU activation followed by another linear transformation. The dimensionality of its input and output is d_{model} .

3.4 Dropout

Dropout is applied to the output of the three sub-layers as well as after the ReLU activation in the third sub-layer. We use a rate of $P_{drop} = 0.1$, as is used in (Vaswani et al., 2017). Moreover, residual connection and layer normalisation are used only in the third sub-layer.

4 Data

Our proposed Sig-Transformer Encoder can be used for general purpose representation learning, and any classifier or regressor can be added on top of it to perform down-stream tasks. We conduct all of our experiments with a medical prescription data for the class of “beta-blockers (ATC code C07)”, provided by the Karolinska University Hospital. It is sampled from a much larger database of 41 million medical prescription records. This dataset contains 3,852 non-duplicated prescriptions written in Swedish. Some examples of annotated prescriptions are shown in Table 1.

The medical practitioners at the hospital have annotated three labels, in which we use QUANTITY for regression while QUANTITY TAG and INDICATION as two classification tasks:

- QUANTITY: the total amount of tablets or capsules prescribed. The values are multiples of 0.5;
- QUANTITY TAG (5 classes): the label or tag of the quantity prescribed to the patient;
 1. *Not Specified*: the quantity was not specified in the prescription;
 2. *Complex* : a range of quantities was given. In that case, the quantity is an average between the minimum and the maximum quantities;
 3. *PRN* : prescription to take only if needed;
 4. *As Per Previous Prescription*: refers to guidance in previous prescription;
 5. *Standard*: standard prescription;
- INDICATION (5 classes): the purpose of the prescription. It originally had 44 classes where many account for only one record⁴. The

⁴The original 44 classes and the number of prescriptions per class, can be found in Figure 5.

medical practitioners have aggregated them into 5 medically meaningful classes: *Cardiac*, *Tremors*, *Migraine*, *Others*, and *NA* (Not Annotated).

One challenge of the data is the remarkable class imbalance in its labels, as shown in Table 2. It’s still prominent in the INDICATION field after aggregating the original classes, where *Cardiac* and *NA* account for majority of the instances, while the other 3 classes are seldom annotated in the data. Another challenge, which is prevalent in electronic health record (EHR) datasets, is the free-form nature of its text and the writing styles vary considerably between different doctors.

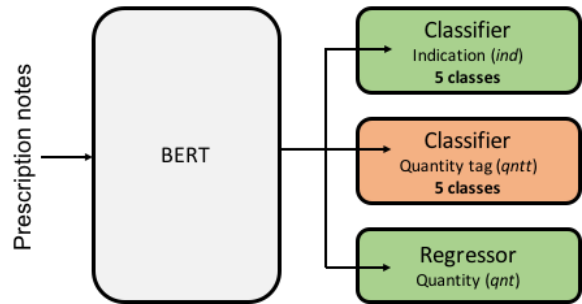


Figure 3: Multi-task learning architecture without Sig-Transformer Encoder (STE).




Figure 4: Our proposed multi-task learning architecture with Sig-Transformer Encoder (STE).

4.1 Data Preprocessing

As mentioned in Section 1 we experiment two different embedding models for the Swedish prescription records, in which one of them, namely ClinicalBERT (Huang et al., 2020), requires its input text to be English. We use the Google Translate API to obtain translated English prescriptions. Minimal text preprocessing was taken, and all prescription text including non-ASCII characters (e.g. å, ä, ö) was read with UTF-8 encoding.

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

Table 1: Example prescriptions with translations and annotations. The second column is the English translation obtained from Google Translate API. The last three columns are the labels of interest which we aim to extract automatically. Some longer example prescriptions can be viewed in Table 8.

QUANTITY TAG			INDICATION		
Standard	3489	90.6%	Cardiac	2980	77.4%
PRN	89	2.3%	Tremors	81	2.1%
APPP	40	1.0%	Migraine	69	1.8%
Complex	29	0.8%	Other	15	0.4%
NS	205	5.3%	NA	707	18.4%

Table 2: Class distribution for QUANTITY TAG and INDICATION. APPP is an abbreviation for *As Per Previous Prescription*, NS stands for *Not Specified* and NA stands for *Not Annotated*.

5 Multi-Task Learning

Our objective is to automatically find the QUANTITY of prescribed medicine as well as QUANTITY TAG and INDICATION of the prescription, for each prescription note. We formulate this as a multi-task learning problem consisting of a regression task and two classification tasks. As depicted in Figure 3, such multi-task learning model uses BERT or any of its variants to encode each input prescription note represented by the respective [CLS] token, then we can add three separate predictors on top of it and all three jointly learn with the embedding layer. We use this architecture as one of our baseline approaches.

In our proposed multi-task learning architecture, we add Sig-Transformer Encoder (STE) in-between the embedding layer (BERT) and the three predictors in order to learn the sequential order informa-

tion in the input text encoded by the token-level BERT representations (instead from [CLS]). The three predictors and STE are jointly trained. We compare the two architectures and other baseline models in Section 6.

6 Experiments

6.1 Experiment Setup

As described in Section 2, BERT and its variants have been used in a range of clinical machine learning tasks. Considering our data is in Swedish we explore two different approaches for encoding the prescription notes: (1), Apply Multilingual BERT (M-BERT) (Devlin et al., 2018) directly to the Swedish text; (2), Translate the prescriptions to English as described in Section 4.1, and then encode the translated text with ClinicalBERT (Huang et al., 2020). ClinicalBERT is fine-tuned on clinical notes, therefore it has embedded more domain knowledge and has proven to be more effective over BERT-base on several clinical tasks. However, any translation error would affect the performance of subsequent models.

We use the PyTorch implementation of Transformers by Hugging Face⁵ for loading M-BERT and ClinicalBERT, and Signatory⁶ (Kidger and Lyons, 2020) for differentiable computations of

⁵<https://github.com/huggingface/transformers>

⁶<https://github.com/patrick-kidger/signatory>

signatures on GPU. We use a linear layer as the classifier for QUANTITY TAG, and two-layer network for QUANTITY and INDICATION respectively. More details of the implementation choices are described in Appendix A.1. We use cross-entropy (CE) loss for the two classification tasks and mean-square error (MSE) for the regression task. The overall loss function is a weighted average of the loss functions of all three tasks:

$$\mathcal{L} = \alpha_{qnt} \times MSE_{qnt} + \beta_{qntt} \times CE_{qntt} + \beta_{ind} \times CE_{ind}$$

where qnt stands for quantity, $qntt$ stands for quantity tag, ind stands for indication, and $\alpha_{qnt} + \beta_{qntt} + \beta_{ind} = 1$. α_{qnt} , β_{qntt} and β_{ind} are parameters to be optimised.

To compare model performance we use the simple multi-task learning models described in Section 5 without adding STE (Figure 3), as the baseline systems. We also replace STE with LSTM for comparison as both learn sequential information from the input data⁷. All experiments are conducted using 5-fold cross validation, where we use 3 folds for training, 1 fold for validation and 1 fold as the test set at each iteration. We repeat the experiments for each model with 10 different random seeds and average the scores as the final result.

6.2 Results

The results for all three tasks are summarised in Table 3. We first notice all models have failed to recognise the INDICATION classes, which we will discuss in Section 6.3. In this section we mainly discuss the performance for QUANTITY and QUANTITY TAG.

Among the three models without the use of LSTM or STE, both ClinicalBERT and Multilingual BERT (M-BERT) outperform the BERT-base model in all three tasks. As for the three LSTM-based models, using the BERT-base representations achieves the best performance while ClinicalBERT and M-BERT are comparable with each other. Among our proposed STE-based models, M-BERT+STE outperforms the other two variants.

Comparing the results across the board, adding LSTM or STE improves the model performance in most cases except for ClinicalBERT, where the

⁷Unlike the Transformer model, recurrent neural networks (RNNs) such as LSTM, model relative and absolute positions along the time dimension directly through their sequential structure (Shaw et al., 2018).

addition of LSTM or STE has worsened the performance for ClinicalBERT. We think the potential discrepancies in the translation of the Swedish prescriptions have possibly affected the working of LSTM and STE, as both models take all the translated token embeddings of each prescription as input where the simple ClinicalBERT baseline only takes the encoding of the [CLS] token as input. Comparing between the two different prescription embedding approaches, translating them then apply ClinicalBERT clearly beats using the multilingual-BERT when LSTM or STE is not added. When we incorporate LSTM or STE in the model, the multilingual embedding approach has obtained better results in most cases. Overall, one of our proposed models, M-BERT+STE, achieves the best performance for the QUANTITY and QUANTITY TAG tasks⁸.

In Table 6 we present results obtained by each model across different classes of QUANTITY TAG. We can see the *Standard* class (i.e. standard prescription) followed by *PRN* (i.e. prescription to take only if needed) and *APPP* (i.e. as per previous prescription) have received better classifications by the three STE models. The models also performed well in the *Not Specified* class. The *Complex* class refers to a range of quantities given in the prescription. As expected, many models do not perform well in this class. The overall best performing model, M-BERT+STE, also performs consistently well across different classes of QUANTITY TAG.

6.3 The difficult case of Indication

The class imbalance is remarkable for QUANTITY TAG and INDICATION in this dataset, as described in Section 4. The strong results obtained from many models for QUANTITY TAG suggests albeit the class imbalance the boundaries between its classes are distinguishable and can be learnt from fewer examples. INDICATION is a more challenging task. The aggregation of the original 44 classes to the final 5 classes has reduced the severity of the class imbalance issue⁹. However, the new classification becomes much less distinguishable and the new classes are less distinct as each class now contains

⁸The results obtained by M-BERT+STE and ClinicalBERT are different at the significance level of 0.1, with $p = 0.068$ and 0.065 for QUANTITY and QUANTITY TAG respectively, using Wilcoxon signed-rank test.

⁹We have experimented with the original 44 classes, and obtained terrible results. We have also tried setting class weights to alleviate the class imbalance issue with no significant difference observed in the final results.

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

Table 3: Performance comparison between our proposed STE-based approaches and baseline models. The QUANTITY task is measured in mean squared error (MSE), while for QUANTITY TAG and INDICATION we use Macro F-1 score. Base refers the original BERT-base model, M-BERT is the Multilingual BERT model pretrained on Swedish text data. STE refers to Sig-Transformer Encoder.

a range of topics, as shown in Figure 5. Such low scores are also observed across different classes of INDICATION in Table 7.

6.4 Ablation Study

The positional encoding proposed in (Vaswani et al., 2017) is based on sinusoids of varying frequency, and the authors hypothesise this would help the model to generalise to sequences of variable lengths during training. Our proposed architecture has both the absolute positional representations of words and the signature transform that is also invariant to the sequence length for encoding order of events. Here we conduct ablation study by removing the positional encoding or the signature transform components from the best performing model, M-BERT+STE, and compare the results. As shown in Table 4, we obtain worse performance without positional encoding or without signature transform, for all three tasks, suggesting the contributions of using the absolute positional encoding and the proposed integration of signature transform.

Model	Quantity	Quantity Tag	Indication
Full model	0.15	0.92	0.05
w/o PE	0.38	0.81	0.01
w/o ST	0.31	0.77	0.02

Table 4: Ablation study of the best performing M-BERT + STE model, removing positional encoding (PE) or signature transform (ST).

7 Conclusions and Future Work

In this work, we propose a new extension to the Transformer architecture, named Sig-Transformer

Encoder (STE), by incorporating signature transform with the self-attention mechanism. As a preliminary study, we aimed to automatically extract information related to *quantity*, *quantity tag* and the *indication* label from a Swedish medical prescription dataset. We demonstrated good performance in two of the three tasks in a multi-task learning framework. Lastly, we compared two embedding approaches between applying ClinicalBERT (on translated text) and Multilingual BERT (M-BERT). Although one of our proposed models, namely M-BERT+STE, reported the best performance for *quantity* and *quantity tag*, all models failed to perform on the *indication* task which we provided possible explanations.

For future work, we plan to apply the proposed STE models to the much larger prescription database and investigate ways to improve the labelling for *indication*. We also plan to benchmark our models on other clinical tasks such as predicting hospital readmission (Huang et al., 2020), as well as the more general NLP tasks evaluated in (Devlin et al., 2018). Moreover, a further investigation to better understand the contribution of signature transform and the interplay between signature and attentions would be very insightful.

Acknowledgments

This work was supported by the MRC Mental Health Data Pathfinder award to the University of Oxford [MC_PC_17215], by the NIHR Oxford Health Biomedical Research Centre and by the The Alan Turing Institute under the EPSRC grant EP/N510129/1. We would like to thank Dr. Remus-Giulio Anghel and Dr. Yasmina Molero Plaza for

providing the data and sharing information about the data.

References

- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical bert embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78.
- Imanol Perez Arribas, Guy M Goodwin, John R Geddes, Terry Lyons, and Kate EA Saunders. 2018. A signature-based machine learning model for distinguishing bipolar disorder and borderline personality disorder. *Translational psychiatry*, 8(1):1–7.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Juan M Banda, Martin Seneviratne, Tina Hernandez-Boussard, and Nigam H Shah. 2018. Advances in electronic phenotyping: from rule-based definitions to machine learning models. *Annual review of biomedical data science*, 1:53–68.
- Ilya Chevyrev and Andrey Kormilitzin. 2016. A primer on the signature method in machine learning. *arXiv preprint arXiv:1603.03788*.
- Thomas Davenport and Ravi Kalakota. 2019. The potential for artificial intelligence in healthcare. *Future healthcare journal*, 6(2):94.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. Domain-specific language model pretraining for biomedical natural language processing. *arXiv preprint arXiv:2007.15779*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. 2020. Clinicalbert: Modeling clinical notes and predicting hospital readmission. In *Proc. ACM Conference on Health, Inference, and Learning (CHIL)*.
- 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(1):1–9.
- Katikapalli Subramanyam Kalyan and S Sangeetha. 2020. SecNlp: A survey of embeddings in clinical natural language processing. *Journal of biomedical informatics*, 101:103323.
- 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:2001.00706*.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Terry Lyons. 2014. Rough paths, signatures and the modelling of functions on streams. In *Proceedings of the International Congress of Mathematicians*, pages 163–184.
- Terry J Lyons. 1998. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310.
- James H Morrill, Andrey Kormilitzin, Alejo J Nevado-Holgado, Sumanth Swaminathan, Samuel D Howison, and Terry J Lyons. 2020. Utilization of the signature method to identify the early onset of sepsis from multivariate physiological time series in critical care monitoring. *Critical Care Medicine*, 48(10):e976–e981.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. 2019. Transfer learning in biomedical natural language processing: An evaluation of bert and elmo on ten benchmarking datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 58–65.
- Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468.
- Yuqi Si, Jingqi Wang, Hua Xu, and Kirk Roberts. 2019. Enhancing clinical concept extraction with contextual embeddings. *Journal of the American Medical Informatics Association*, 26(11):1297–1304.
- Csaba Toth and Harald Oberhauser. 2020. Bayesian learning from sequential data using gaussian processes with signature covariances. In *Proceedings of the International Conference on Machine Learning (ICML)*.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Bo Wang, Maria Liakata, Hao Ni, Terry Lyons, Alejo J Nevado-Holgado, and Kate Saunders. 2019. A path signature approach for speech emotion recognition. In *Interspeech 2019*, pages 1661–1665. ISCA.
- Bo Wang, Yue Wu, Niall Taylor, Terry Lyons, Maria Liakata, Alejo J Nevado-Holgado, and Kate EA Saunders. 2020. Learning to detect bipolar disorder and borderline personality disorder with language and speech in non-clinical interviews. In *Interspeech 2020*. ISCA.
- Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. 2018. Clinical information extraction applications: a literature review. *Journal of biomedical informatics*, 77:34–49.
- Zecheng Xie, Zenghui Sun, Lianwen Jin, Hao Ni, and Terry Lyons. 2018. 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, Lianwen Jin, and Manfei Liu. 2016. Deepwriterid: An end-to-end online text-independent writer identification system. *IEEE Intelligent Systems*, 31(2):45–53.
- 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*.

A Appendices

A.1 Implementation Details

Our classifier for INDICATION is a two-layer network, in which the first layer has the size of 50 and 5 for the second layer. ReLU activation and dropout are used in-between the two layers. The dropout rate is set to be 0.1. We also use a two-layer network for the QUANTITY regression task, in which the first layer has the size of 10 and its second layer predicts the value for *quantity*. We also add a dropout the rate 0.1 in-between the two layers. Our classifier for QUANTITY TAG has only one layer the size of 5.

We reduce the dimension of attended vector before signature transform from 768 to 32, as depicted in Figure 2 (left). We set the order of truncated signature to be 2, and employ 8 parallel attention layers (or heads). We ran grid-search over two choices of learning rate lr : $[0.00003, 0.00005]$, and 10 randomly sampled triplets of α_{qnt} , β_{qntt} and β_{ind} . Each value of α_{qnt} , β_{qntt} and β_{ind} is between 0 and 1 respectively, while we make sure each combination of the three values sums to 1.

A.2 Signature of paths

We begin with the definition of the signature, using more traditional notation of stochastic calculus.

Definition 1. Let $X = (X^1, \dots, X^d)$ be a path in \mathbb{R}^d . The signature of X is defined as the infinite collection of iterated integrals:

$$\begin{aligned} S(X) &= \left(\int \dots \int_{a < t_1 < \dots < t_k < b} dX_{t_1} \otimes \dots \otimes dX_{t_k} \right)_{k \geq 0} \\ &= \left(\left(\int \dots \int_{a < t_1 < \dots < t_k < b} dX_{t_1}^{i_1} \dots dX_{t_k}^{i_k} \right)_{1 \leq i_1, \dots, i_k \leq d} \right)_{k \geq 0} \end{aligned}$$

where $dX_t = \frac{dX_t}{dt} dt$ and the $k = 0$ term is taken to be $1 \in \mathbb{R}$.

Definition 2. The truncated signature of order N of X is defined as:

$$S^N(X) = \left(\int \dots \int_{a < t_1 < \dots < t_k < b} dX_{t_1} \otimes \dots \otimes dX_{t_k} \right)_{0 \leq k \leq N}.$$

The dimension of the truncated signature explodes exponentially with the input path dimension:

Proposition 1. For any $d \geq 1$, the truncated signature of order N of a d -dimensional path has the dimension of:

$$\sum_{k=0}^N d^k = \frac{d^{N+1} - 1}{d - 1}$$

In practice, the term of order 0 is dropped as it is always equal to 1. For clarity, we define:

$$S(X)^{i_1, \dots, i_k} = \int \dots \int_{a < t_1 < \dots < t_k < b} dX_{t_1}^{i_1} \dots dX_{t_k}^{i_k}$$

with $1 \leq i_1, \dots, i_k \leq d$, so that:

$$\begin{aligned} S(X) &= \left((S(X)^{i_1, \dots, i_k})_{1 \leq i_1, \dots, i_k \leq d} \right)_{k \geq 0} \\ &= (1, S(X)^1, \dots, S(X)^d, S(X)^{1,1}, S(X)^{1,2}, \dots). \end{aligned}$$

d_{presig}	$order_{sig}$	d_{sig}	d_{presig}	$order_{sig}$	d_{sig}
512	1	512	16	3	4K
512	2	262K	8	2	72
256	2	66K	8	3	584
128	2	16K	8	4	5K
128	3	2M	4	4	340
64	2	4K	4	5	1365
64	3	266K	4	6	5K
32	2	1057	2	9	1022
32	3	34K	2	10	2K
16	2	272	2	12	8K

Table 5: The number of dimensions (d_{sig}) of the truncated signature is determined by the size of its input (d_{presig}) and the order of truncation selected ($order_{sig}$).

Model	QUANTITY TAG					QUANTITY
	Standard	APPP	PRN	Complex	NS	
Base	0.71	0.50	0.10	0.76	0.95	0.50
ClinicalBERT	0.83	0.97	0.89	0.99	0.79	0.21
M-BERT	0.94	0.41	0.04	0.81	0.98	0.41
Base + LSTM	0.93	0.36	0.99	0.77	0.74	0.45
ClinicalBERT + LSTM	0.90	0.23	0.99	0.61	0.49	0.47
M-BERT + LSTM	0.29	0.99	0.72	0.47	0.93	0.50
Base + STE	0.97	0.84	0.77	0.44	0.88	0.23
ClinicalBERT + STE	0.99	0.70	0.85	0.65	0.65	0.36
M-BERT + STE	0.86	0.97	1.00	0.89	0.86	0.15

Table 6: Model performance comparison for QUANTITY and also across different classes in QUANTITY TAG. APPP: As Per Previous Prescription; NS: Not Specified.

Model	INDICATION				
	Cardiac	Tremors	Migraine	Others	NA
Base	0.03	0.05	0.02	0.27	0.03
ClinicalBERT	0.00	0.24	0.03	0.04	0.13
M-BERT	0.05	0.01	0.09	0.35	0.00
Base + LSTM	0.02	0.00	0.00	0.31	0.00
ClinicalBERT + LSTM	0.01	0.03	0.34	0.08	0.02
M-BERT + LSTM	0.05	0.00	0.00	0.04	0.03
Base + STE	0.00	0.00	0.02	0.08	0.05
ClinicalBERT + STE	0.00	0.26	0.05	0.02	0.00
M-BERT + STE	0.00	0.11	0.00	0.00	0.15

Table 7: Model performance comparison across different classes in INDICATION. NA: Not Annotated.

Indication (44)	
CARDIAC - HYPERTENSION - PALPITATIONS	49
CARDIAC - PALPITATIONS - TREMORS	1
CARDIAC - PALPITATIONS - ANGINA	1
CARDIAC - PALPITATIONS - ATRIAL FIBRILLATION	1
CARDIAC - ATRIAL FIBRILLATION	57
CARDIAC - DYSRHYTHMIA	101
CARDIAC - HEART FAILURE	60
CARDIAC - HYPERTENSION - ANGINA - PALPITATIONS	1
CARDIAC - ATRIAL FIBRILLATION - DYSRHYTHMIA/RATE	1
CARDIAC - HYPERTENSION - DYSRHYTHMIA	19
CARDIAC - MIGRAINE	1
CARDIAC - DYSRHYTHMIA/RATE	114
CARDIAC - HYPERTENSION - ESOPHAGEAL VARICES	1
CARDIAC - PALPITATIONS	162
CARDIAC - HYPERTENSION - HEART FAILURE	4
CARDIAC - HYPERTENSION - MIGRAINE	12
CARDIAC - HYPERTENSION - ANGINA	28
CARDIAC - TREMORS	2
CARDIAC - DYSRHYTHMIA - MIGRAINE	1
CARDIAC - ANGINA - ATRIAL FIBRILLATION	1
CARDIAC - HEPATIC CIRCULATION	1
CARDIAC - ANGINA	34
CARDIAC - DYSRHYTHMIA/RATE - PALPITATIONS	1
CARDIAC - HYPERTENSION	1655
CARDIAC - HYPERTENSION - ATRIAL FIBRILLATION	2
CARDIAC - ESOPHAGEAL VARICES	2
CARDIAC	608
CARDIAC - ANGINA - DYSRHYTHMIA/RATE	1
CARDIAC - HYPERTENSION - DYSRHYTHMIA/RATE	50
CARDIAC - PALPITATIONS - MIGRAINE	1
CARDIAC - HYPERTENSION - TREMORS	8
ESOPHAGEAL VARICES	1
ANXIETY - CARDIAC - PALPITATIONS	3
ANXIETY	6
ANXIETY - MIGRAINE	1
HYPERTHYROIDISM	1
HYPERGLYCEMIA	2
CHEST PAIN	1
TREMORS	77
TREMORS - CARDIAC - HYPERTENSION	1
TREMORS - ANXIETY	1
TREMORS - CARDIAC - PALPITATIONS	2
MIGRAINE	69
NA	707

Figure 5: Number of prescriptions per indication class, where the indication label has the original 44 classes. *NA* stands for *Not Annotated*.

Swedish	Translated English	Indication	Quantity	Quantity Tag
0.5 TABLETTER 2 GANGER DAGLIGEN. FOR HJARTRYTMEN OCH SANKER BLODTRYCKET TABLETTERNA SVALJES HELA (KAN DELAS VID SVALJSVARIGHETER, MEN FAR EJ K	0.5 tablets 2 times daily. For the heart rhythm and reduces blood pressure The tablets are swallowed whole (can be divided at swallowing durations, but don't k	Cardiac- hypertension	1	Standard
1 TABLETT 1 GANG DAGLIGEN I 3 VECKOR. EVENTUELL HOJNING TILL 2 TABLETTER DAGLIGEN BEROENDE PA KONTROLLEN OCH SYMPTOMLINDRING, MOT HJARTKLAPPNING	1 table 1 time daily for 3 weeks. Successful increase for 2 tablets daily dependent on control and symptoms relief, against palpitations	Cardiac- palpitations	1	Complex
1 TABLETT 1 GANG DAGLIGEN. EFTERHAND EVENTUELL OKNING TILL EN TABLETT MORGON OCH LUNCH MEN BORJA MED EN TABLETT TIDIG MORGON, MOT SKAKNINGAR.	1 tbele 1 time daily. Previously opening to a tablet morning and lunch but begin with a tablet early morning, against shakes	NA	1	Standard
1/2 TABLETT PA MORGONEN , 1/2 TABLETT PA LUNCHEN OCH 1/2 TABLETT PA KVALLEN MOT TREMOR, BLODTRYCKSREGLERANDE	1/2 table on morning, 1/2 table on lunch and 1/2 table on evening against tremor, blood pressure control	NA	1.5	Standard
1 TABLETT 1 GANG DAGLIGEN MOT HOGT BLODTRYCK DU BOR BESTALLA LAKARTIDLAMPLIGEN MAJ JUNI FOR KONTROLL BT DIABETES.	1 tablet 1 time daily against high blood pressure you should order the painting lighting may june for control BT diabetes	Cardiac- hypertension	1	Standard
2 depottabletter kl. 08, 1 depottablett kl. 20. Dagligen. Mot hOgt blodtryck	2 prolonged-release tablets at. 08, 1 prolonged-release tablet at. 20. Daily. Against high blood pressure	Cardiac- hypertension	3	Standard

Table 8: Longer example prescriptions with translations and annotations.