King's Research Portal

Document Version
Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Wu, H., Ibrahim, Z., Iqbal, E., & Dobson, R. J. B. (Accepted/In press). Encoding Medication Episodes for Adverse Drug Event Prediction. In *The Thirty-Sixth Annual International Conference of the British Computer Society's Specialist Group on Artificial Intelligence*

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Encoding Medication Episodes for Adverse Drug Event Prediction

Honghan Wu and Zina M. Ibrahim and Ehtesham Iqbal and Richard JB Dobson

Abstract Understanding the interplay among the multiple factors leading to Adverse Drug Reactions (ADRs) is crucial to increasing drug effectiveness, individualising drug therapy and reducing incurred cost. In this paper, we propose a flexible encoding mechanism that can effectively capture the dynamics of multiple medication episodes of a patient at any given time. We enrich the encoding with a drug ontology and patient demographics data and use it as a base for an ADR prediction model. We evaluate the resulting predictive approach under different settings using real anonymised patient data obtained from the EHR of the South London and Maudsley (SLaM), the largest mental health provider in Europe. Using the profiles of 38,000 mental health patients, we identified 240,000 affirmative mentions of dry mouth, constipation and enuresis and 44,000 negative ones. Our approach achieved 93% prediction accuracy and 93% F-Measure.

1 Introduction

Adversities associated with prescribed medication can seriously affect the patient's wellbeing [1] and present a real financial burden on healthcare providers (estimated to lead to an annual cost of 466 million pounds in the United Kingdom alone [3]). The problem arises from the fact that trial-tested prescribed drugs are not evaluated for long-term effectiveness in diverse (and possibly comorbid) patient populations, resulting in many unknowns with respect to the possibilities of the onset of adverse drug reactions (ADRs). Currently, discovering the factors affecting patients' response to treatment relies on spontaneous reporting systems that rely on patient and clinician data entry and resulting highly under-reported ADR instances [5]. However, there is a large body of unutilised knowledge embedded in the Elec-

Honghan Wu, Zina M. Ibrahim, Ehtesham Iqbal and Richard JB Dobson
King's College London, London, United Kingdom, e-mail: {honghan.wu, zina.ibrahim, ehtesham.iqbal, richard.j.dobson}@kcl.ac.uk

tronic Health Records (EHRs) of hospitals, containing valuable information regarding treatment responses as well as patient profiles and disease trajectories.

Analysing the factors associated with ADRs using EHRs resolves to the task of uncovering the temporal associations connecting the different factors. Things are more complicated in long-term debilitating and comorbidity illnesses whereby ADRs are associated with multiple, overlapping and long medication episodes. From a technical perspective, a challenging task lies in modelling the time series of medication episodes in an effective way so that they can be fed with other influencing factors into predictive algorithms.

In this paper, we propose a bitwise encoding mechanism that can capture the temporal precedence and duration information of medication episodes as well as their distance from the time of inspection. We present an initial study of using the resulting model to predict the onset of three ADRs, mainly dry mouth, constipation and enuresis, from a large mental health register.

2 Temporal Encoding for Patient's Medication Episodes

Here, we present the novel vector-based encoding of dynamic and multiple medication episodes. Figure 1 illustrates an example encoding scenario. The upper part shows the sample medication episodes of a patient that are related to a certain inspecting time spot (marked as AE date in the figure). The encoding is realised through the following steps.

Fig. 1 The bitwise encoding of dynamic medication episodes for a given adverse event

1. From the the inspecting time spot, look back to a certain period of time, e.g., 20 days. In our experiments, we selected 30 days.
2. Split the time period into intervals using a unit of time, e.g., a day or a week. Number each interval in an ascendent order (started from zero) starting from the

inspecting time spot. For example, 4 intervals are identified and numbered in Figure 1

3. Add the patient’s medication episodes to the timeline.
4. For each medication episode, allocate a sequence of bits aligned with the interval order we obtained in step 2. The interval numbered zero is aligned with the most significant bit in the sequence. For each interval, set the corresponding bit to 1 if the interval intersects with the episode, and to 0 otherwise. For example, Medication Episode 2 (M2) in Figure 1 spans across interval number 1 and 2. Its encoding is 0110 in binary code.
5. Repeat step 4 for all medication episodes. We will get a vector representing all the relevant medication episodes of the inspecting time spot.

2.1 Predictive Model

Having illustrated the encoding of the temporal information embedded within medication episodes, we now show how this information can be combined with additional background knowledge as well as external ontologies into a feature vector for prediction.

Our predictive model considers two sets of features: 1) demographics, 2) the generated dynamic medication episodes which are further enriched by an ontology. Equation 1 shows the features of demographic information of a patient p at a given time t where $age@t$ means the patient’s age at the inspecting time t .

$$F_d(p, t) = (age@t, gender, ethnicity) \quad (1)$$

Equation 2 gives the feature calculation from dynamic medication episodes given a patient p , an adverse event type a and a time t .

$$F_m(p, a, t) = (\alpha \cdot \text{diag}(\mathbf{S}_a) + \text{diag}(\mathbf{1})) \cdot \mathbf{M}_{p,t} \quad (2)$$

$\mathbf{S}_a$ is a binary vector describing drugs’ relation with side effect a . For example, suppose there are two drugs (d_1, d_2), and d_1 has the side effect a but d_2 does not. Then, the $\mathbf{S}_a$ is $[1, 0]$. Such vector is derived from the drug ontology. $\mathbf{M}_{p,t}$ is the vector generated from the medication episodes of the patient p at a given time t using our encoding method. α is the weighting factor for adjusting side effect knowledge importance in the final prediction model.

The final feature vector for prediction is a concatenation of the above two as shown in equation 3. In this paper, we view the ADR prediction as a classification problem. Therefore, the feature vector is used as the inputs for training and testing classification algorithms on our datasets. Various classifiers have been tested. The experiment section will give the detail about model selection.

$$F(p, a, t) = (F_m(p, a, t), F_d(p, t)) \quad (3)$$

3 Implementation and Evaluation

Data Source and Preparation

We used data extracted from the Clinical Record Interactive Search System (CRIS) [4] to evaluate our encoding. CRIS is a database containing a de-identified replica of the EHRs used by the South London and Maudsley Foundation Trust (SLaM), the largest mental health provider in Europe. SLaM serves over 1.2 million patients and stores much of its clinical records and prescribing information in unstructured free text format. The ADR and drug episode information were extracted using an in-house developed natural language processing tools in conjunction with manual annotation [2]

In this preliminary study, we focus on three types of adverse events identified in the CRIS registry: dry mouth (#event 58,347), constipation (#event 86,602) and enuresis (#93,366), which involves 20,795 distinct patients with mental health disorders. For each identified AE, we pick up the patient’s past 30 days medication episodes. It is worth mentioning that a large proportion ($\geq 62\%$ on average) of AEs are associated with multiple medications. To generate negative data items (none adverse event), we picked 18,038 patients who have medication episode data in CRIS registry but never had the three types of adverse events reported.

The dataset was split into a proportion of 80/20 for training and testing. 10-folds cross validation has also been used. The former setting achieved better performance across three AE types.

Demographic vs Medication Episode Features

We have two types of feature sets - demographic features and medication episode ones. Comparing their performances will reveal some insights about what factors are more likely to be associated with adverse events. When only using demographic features, the F-Measure is around 70%. But for AE class of dry mouth type, the performance is extremely low - 12% True Positive and 20% F-Measure. Table 1 gives the performance of dynamic medication episode features. In all cases, about 90% F-Measure has been achieved with the lowest at 85% at AE class of dry mouth and the highest at 96% at none AE class of the same case. In summary, it is quite obvious that dynamic medication episode features are much better indicators for predicting Adverse Events. The combination of the two sets of features can achieve the best results as shown in Table 2. Both average Precision and F-Measure are around 93%.

Ontology-based Feature Dimension Reduction

The semantics in the drug ontologies can be utilised to combine semantically similar drugs so that the dimensions of the feature vector F_m (Equation 2) can be reduced

Table 1 Medication Episode-Feature Only Results

ADE Type	TP Rate	FP Rate	Precision	Recall	F-Measure	ROC Area	Class
Dry Mouth	0.83	0.04	0.86	0.83	0.85	0.97	AE
	0.96	0.17	0.95	0.96	0.96	0.97	N-AE
Weighted Avg.	0.93	0.14	0.93	0.93	0.93	0.97	
Constipation	0.90	0.07	0.92	0.90	0.91	0.97	AE
	0.93	0.10	0.91	0.93	0.92	0.97	N-AE
Weighted Avg.	0.92	0.08	0.92	0.92	0.92	0.97	
Enuresis	0.88	0.07	0.93	0.88	0.90	0.96	AE
	0.93	0.12	0.88	0.93	0.90	0.96	N-AE
Weighted Avg.	0.90	0.10	0.90	0.90	0.90	0.96	

Table 2 Medication Episode and Demographic Combined Results

ADE Type	TP Rate	FP Rate	Precision	Recall	F-Measure	ROC Area	Class
Dry Mouth	0.89	0.05	0.85	0.89	0.87	0.98	AE
	0.95	0.11	0.97	0.95	0.96	0.98	N-AE
Weighted Avg.	0.94	0.10	0.94	0.94	0.94	0.98	
Constipation	0.93	0.08	0.91	0.93	0.92	0.97	AE
	0.92	0.08	0.93	0.92	0.93	0.97	N-AE
Weighted Avg.	0.92	0.08	0.92	0.92	0.92	0.97	
Enuresis	0.92	0.09	0.92	0.92	0.92	0.96	AE
	0.91	0.08	0.91	0.91	0.91	0.96	N-AE
Weighted Avg.	0.92	0.08	0.92	0.918	0.92	0.96

Fig. 2 Feature Dimension Reduction by Drug Ontology

significantly (from 226 for all drugs to 47 for using direct parent category). The performance differences between these two settings are illustrated in Figure 2. For dry mouth and constipation cases, the performances (average classification accuracy) decreased in about 1 per cent in average when using drug classes, while the accuracy did improve in the enuresis case when the feature dimension was reduced.

This is a very interesting observation worth further investigation to see whether it is a single special case or the drugs related to enuresis make the case special. Figure 2 also reveals that using less feature dimension can significantly increase the training speed, which is not surprising.

4 Conclusion

We reported a retrospective study of predicting adverse events in EHRs from South London and Maudsley Foundation Trust (SLaM), the largest mental health provider in Europe. The prediction model puts a special focus on effective approaches of modelling the dynamic and multiple medication episodes, which are very common among patients with mental health disorders (observed among $\geq 62\%$ events in our data). Specifically, a novel bitwise encoding approach is introduced for capturing the medication trajectories, which has proved to be more effective (10% accuracy improvement) than static medication modelling settings.

Acknowledgement This work has received support from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644753 (KConnect) and National Institute for Health Research (NIHR) Biomedical Research Centre and Dementia Unit at South London and Maudsley NHS Foundation Trust and Kings College London.

References

1. Edwards, I., Aronson, J.: Adverse drug reactions: definitions, diagnosis, and management. *The Lancet* **356**(9237), 12551259 (2000)
2. Iqbal, E., Mallah, R., Jackson, R.G., Ball, M., Ibrahim, Z.M., Broadbent, M., Dzahini, O., Stewart, R., Johnston, C., Dobson, R.J.B.: Identification of adverse drug events from free text electronic patient records and information in a large mental health case register. *PloS one* **10**(8), e0134,208 (2015)
3. Pirmohamed, M., James, S., Meakin, S., Green, C., Scott, A.K., Walley, T.J., Farrar, K., Park, B.K., Breckenridge, A.M.: Adverse drug reactions as cause of admission to hospital: prospective analysis of 18 820 patients. *Bmj* **329**(7456), 15–19 (2004)
4. Stewart, R., Soremekun, M., Perera, G., Broadbent, M., Callard, F., Denis, M., Hotopf, M., Thornicroft, G., Lovestone, S.: The South London and Maudsley NHS foundation trust biomedical research centre (SLaM BRC) case register: development and descriptive data. *BMC psychiatry* **9**(1), 1 (2009)
5. Tatonetti, N.P., Patrick, P.Y., Daneshjou, R., Altman, R.B.: Data-driven prediction of drug effects and interactions. *Science translational medicine* **4**(125), 125ra31–125ra31 (2012)