

Learning Treatment Regimens from Electronic Medical Records

Khanh Hung Hoang and Tu Bao Ho

Japan Advanced Institute of Science and Technology,
1-1 Asahidai, Nomi, Ishikawa, Japan
{hung.hoang,bao}@jaist.ac.jp

Abstract. Appropriate treatment regimens play a vital role in improving patient health status. Although some achievements have been made, few of the recent studies of learning treatment regimens have exploited different kinds of patient information due to the difficulty in adopting heterogeneous data to many data mining methods. Moreover, current studies seem too rigid with fixed intervals of treatment periods corresponding to the varying lengths of hospital stay. To this end, this work proposes a generic data-driven framework which can derive group-treatment regimens from electronic medical records by utilizing a mixed-variate restricted Boltzmann machine and incorporating medical domain knowledge. We conducted experiments on coronary artery disease as a case study. The obtained results show that the framework is promising and capable of assisting physicians in making clinical decisions.

Keywords: treatment regimen, treatment learning, treatment recommendation, electronic medical records

1 Introduction

The two most important issues in healthcare are disease diagnosis and treatment. While many works have been conducted on the problem of diagnosis prediction, the problem of learning treatment regimens has not yet been extensively studied from the research community. This shortage becomes more serious when hospitals essentially need to make efforts to adopt treatment regimens that best fit their available resources. Additionally, it seems hard to have a fixed care plan for a particular disease due to its high dependency on various patient conditions. As a result, capturing treatment regimens in practice turns out to be meaningful for not only assisting physicians in making right clinical decisions but also helping hospitals manage their resources thoroughly.

In principle, treatment regimens could be learned based on the knowledge-driven approach which requires medical domain or expert knowledge. It can be a piece of information written in the literature or accumulated experience gained by physicians during their career. While this approach seems to be reliable, taking various domain knowledge into account is costly and not straightforward in reality. In contrast to the knowledge-driven approach, the data-driven approach

derives treatment patterns from a large number of observations thanks to the availability of electronic medical records in recent years. Studies followed this approach could be found in [2, 5, 6, 8, 9].

Although many interesting results have been achieved, those studies have simply utilized a limited subset of features while many other kinds of data are usually omitted. Such data can be patients' demographics, laboratory test results or clinical notes consisting of signs and symptoms during patients' hospitalization. It is apparent that the more values from those data shared between two patients, the more possibility that the patients are treated with similar regimens. The lack of considering such valuable information simultaneously in current research could be attributed to the poor-feature data used in their experiments. Moreover, even when the above data is made available, it generally exists in form of numerical, binary, categorical, or text format. Such a heterogeneous data is not ready to use for many data mining methods. Another challenge stems from the fact that treatment regimen is typically defined over periods. Each period is distinct from others at milestones where major changes in a patient's health status happen that lead to a notable adjustment in subsequent prescriptions for the patient. Therefore, given a set of prescription records, identifying suitable treatment periods can considerably affect the learned treatment patterns.

This work aims to propose a treatment regimen learning framework which addresses both the above challenges. Our framework first divides patients into clusters from which treatment regimens over periods are discovered then. To overcome the challenges of learning from mixed-type data, we employ a mixed-variate restricted Boltzmann machine (MV.RBM) [11]. The advantage of this model is at its robustness in transforming heterogeneous objects to their homogeneous representations. The new latent representations are in the form of hidden binary vectors that could be further processed easily by clustering methods. To tackle the challenge of treatment period identification, we propose an algorithm which can relatively capture significant changes in prescription indications. Moreover, we also suggest another algorithm which derives treatment regimens from each cluster as a regimen tree. The tree can highlight frequently prescribed drugs and infrequently prescribed drugs inside each patient cluster which would be useful for recommending prescribed drugs to patients.

In short, the main contributions of our work are listed as follows. Firstly, we propose a generic framework which can exploit different kinds of relevant patient records. The framework is superior to others in terms of data utilization. Secondly, we employ both knowledge-driven approach and data-driven approach in our framework. The exploited medical domain knowledge is drug indications and their importance in the treatment for a particular disease. The combination approach used in our framework seems more feasible to deal with the longitudinal property inherent in prescription records. Lastly, we propose a new way to represent treatment regimens flexibly. Frequent drugs are learned from individual level to group level and organized as regimen trees which could be useful for recommending possible regimens to new patients.

2 Related Work

This section provides a brief review of studies about the treatment-related learning problem. Notable works can be found in [2,3,9]. In [3], the authors developed a process mining method to derive clinical pathway from medical behaviors. Their work, however, mainly focused on learning clinical procedures rather than a detailed treatment.

Inspired by the emergence of electronic medical records, recent studies have exploited prescription records which would provide more useful insights about patient treatment. In [2] the authors proposed a probabilistic model that linked patient features and treatment behaviors together to mine treatment patterns. Their model, however, employed many hyperparameters with almost no domain integration. This limitation undermines the model interpretability. Moreover, it was not explicitly described in that work how the chronological order among the learned treatment patterns related. In [9], the inspired work of our research, treatment regimens were derived solely from a set of prescription records. While many typical regimens could be described in an unsupervised mechanism, their prescription-based approach appears to lack of interpretability regarding patient profile and health conditions that lead to the derived regimens. Additionally, although the authors in [9] attempted to describe the chronological order between regimens with predefined treatment periods, their approach capture little medical domain knowledge as well as seems inflexible in dealing with the varying lengths of hospital stay. Regarding the treatment recommendation task, [9] also presented a way to recommend typical treatment regimen for a patient based on demographics and disease severity of patients. This approach, however, seems hard to be applied to new patients whose disease severity may not be recognized at the beginning dates of hospitalization.

3 Methods

In this section, we describe our framework of treatment learning problem. This generic framework is designed for a particular disease. Our approach is based on the assumption that a patient cohort may be divided further into groups of more homogeneous patients who share latent characteristics underlying in patient profile or health status. Patients in one group, therefore, are supposed to be treated by similar care plans that share many parts in common. Fig. 1 illustrates the framework overview. It consists of two main tasks: clustering a cohort of patients and learning treatment regimens for each resulting cluster.

3.1 Data Collection and Preprocessing

Our framework takes medical records of cured patients as trained data. We are interested in the data that characterizes health conditions, for example, demographic information, discharge summary, and laboratory test results. It should be noted that for longitudinal data such as discharge and laboratory indicators,

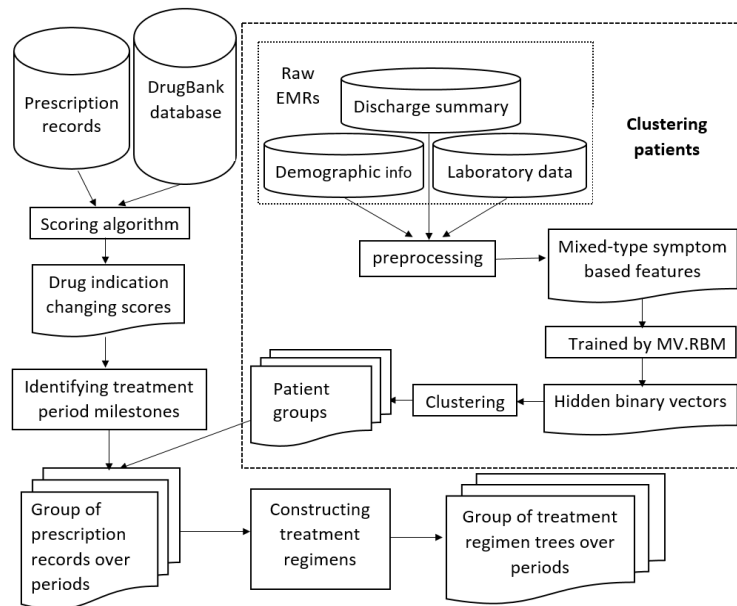


Fig. 1. An overview of the proposed framework for learning treatment regimens.

we only collected the records at the early stage of patients since such longitudinal data is usually not fully available for new patients at the time of admission. This solution is based on the intuition that patients who share initial signs, symptoms and laboratory indicators are likely to be treated in the same way.

After being filtered, patient medical records are encoded as one-hot vectors for categorical data or are normalized to zero-mean unit-variance for numerical data. For discharge summary, only text sections mentioning about the patient history of illness and description about their situation at admission are preferred. We note that segmenting these sections depends on how well-structured discharge summaries were written. In our experiment on MIMIC III database, some clue phrases enabled this solution to become implementable. For simplicity, signs and symptoms mentioned in the segmented text are extracted as new features of the trained patients. Our framework uses the collection processing engine (CPE) component with AggregatePlaintextFastUMLSProcessor provided in cTAKES [7], a well-known tool specifically designed for clinical text processing, to accomplish this task. It is worth noting that extracted signs and symptoms using this tool links to concepts in the Unified Medical Language System (UMLS) [1], the comprehensive ontology built for the biomedical domain.

3.2 Data Representation and Patient Clustering

The encoded data obtained from the previous step contains numerical, binary or categorical values. Such kind of mixed-type data is not ready to fit traditional

clustering methods. Therefore, our framework employs MV.RBM, an extension of the restricted Boltzmann machine for data transformation and representation.

MV.RBM is a RBM where visible nodes are not restricted to binary units. Similar to the original RBM, each binary hidden unit in MV.RBM also tries to capture latent aspects in the inhomogeneous visible units. In other words, MV.RBM could be considered as a model to transform heterogeneous input to homogeneous space. Let $\mathbf{v} = (v_1, v_2, \dots, v_N)$ denote the set of visible features and $\mathbf{h} = (h_1, h_2, \dots, h_K)$ be the set of hidden units. The energy function of MV.RBM is defined more deliberately to handle the mixed variate input.

$$E(\mathbf{v}, \mathbf{h}) = -\left(\sum_i G_i(v_i) + \sum_k b_k h_k + \sum_{ik} H_{ik}(v_i) h_k\right)$$

where $\mathbf{b} = (b_1, b_2, \dots, b_N)$ are biases vectors for hidden layer, $G_i(v_i)$ and $H_{ik}(v_i)$ are specified-type functions. By exploiting the conditional independence property within nodes in a layer of bipartite structure, we can get the following factorization equations: $P(\mathbf{v}|\mathbf{h}) = \prod_{i=1}^N P(v_i|\mathbf{h})$, $P(\mathbf{h}|\mathbf{v}) = \prod_{k=1}^K P(h_k|\mathbf{v})$.

The functions $G_i(v_i)$, $H_{ik}(v_i)$ and corresponding $P_i(v_i|\mathbf{h})$ for each kind of data are given as follows [11].

	$G_i(v_i)$	$H_{ik}(v_i)$	$P_i(v_i \mathbf{h})$
Binary	$a_i v_i$	$w_{ik} v_i$	$\frac{\exp(a_i v_i + \sum_k w_{ik} h_k v_i)}{1 + \exp(a_i + \sum_k w_{ik} h_k)}$
Gaussian	$-v_i^2/2\sigma^2 + a_i v_i$	$w_{ik} v_i$	$\mathcal{N}(\sigma_i^2(a_i + \sum_k w_{ik} h_k), \sigma_i)$
Categorical	$\sum_m a_{im} \delta_m[v_i]$	$\sum_{m,k} a_{imk} \delta_m[v_i]$	$\frac{\exp(\sum_m a_{im} \delta_m[v_i] + \sum_{m,k} w_{imk} \delta_m[v_i] h_k)}{\sum_l \exp(a_{il} + \sum_k w_{ilk} h_k)}$

where a_i , a_{im} are input bias parameters, w_{ik} , w_{imk} are input-hidden weighting parameters. Those with extra subscript m are dedicated for categorical features.




Fig. 2. A MV.RBM for patient records. The green, blue and orange circles represent for binary, categorical and continuous input units. The circles with labels D , S , L indicate demographic, signs/symptoms and laboratory data, respectively.

In our work, we assume features in the preprocessed data are mutually independent given their latent factors. Fig. 2 illustrates our idea to utilize a MV.RBM. We suppose demographic data could receive numerical, binary or categorical values while extracted signs and symptoms are represented as one-hot vectors. Indicator values are assumed to take numerical values. Once the MV.RBM model has been learned, the computable hidden posteriors and hidden states are extracted as transforming features for input \mathbf{v} . Those latent vectors could be used as input of well-known clustering algorithms. In this concrete work, we utilize the learned binary hidden vectors and select the hierarchical clustering algorithm to divide patients into groups. We use the Hamming distance as similarity measurement for binary vectors and the complete linkage which was reported to give low error rate for symmetric distance measurement [10].

3.3 Treatment Period Identification

For each resulting patient cluster obtained from the previous step, prescription records from its patients are collected to derive typical treatment regimens over periods. We represent every drug dr in prescription of patient p as a tuple $dr^p = (name, startdate, enddate, dosage)$ that describes drug name, starting date, ending date of usage and its dosage. Let $\Theta^p = \{dr^p\}$ be the set of drugs given to the patient, and $T^p = \{dr^p.startdate\}$ be the ordered set of dates the patient p was prescribed. As $|T^p|$ varies according to p , we propose an algorithm to split each T^p into the same number of treatment periods. The idea is for each timestamp in T^p , we compute an accumulated score that captures the changes in drug indications that have been delivered to the patient so far. We observe the plot of these scores for many patients in the clusters and decide an appropriate number of periods. The splitting dates for each period are the dates with significant changes in their associated scores.

It is worth noting that in our framework the scoring function takes into account newly prescribed drugs, re-prescribed drugs being stopped using for a while, recently stopped using drugs, or re-prescribed drug with changes in dosage. The aggregate score also gives different weights to those drugs based on their indication. Given a disease de and a set of its common symptoms $Symp^{de}$, we extract from DrugBank database [12] the drugs whose indication description directly mentions about de . We name those drugs as main drugs. Drugs with indication mentioned in $Symp$ are also extracted as symptom-healing drugs. Prescribed drugs for the patient p therefore are classified as main drugs, symptom-healing drugs, and unclassified drugs. The weight of each kind of drug is assigned decreasingly according to its importance for the treatment of de . We denote MDB , SDB as sets of main and symptom-healing drugs which are extracted from DrugBank; w_{main} , w_{symp} , w_{unk} as the weight for main drugs, symptom-healing drugs, and unclassified drugs, respectively. The detailed algorithm for scoring changes in prescribed drug indications for a patient p is presented in Algorithm 1. For readability, we remove the superscript p and use Set notations in the pseudocode.

Algorithm 1: Scoring prescription records

Data: Θ, T, MDB, SDB
Result: return *scores* as a list of accumulated scores
Initialize U as an empty set ; ▷ set of recently delivered drugs
Initialize *scores* as an empty list ;
 $aScore := 0$; ▷ the accumulated score
for each $d \in T$ **do**
 $D := \{dr \mid \forall dr \in \Theta \wedge dr.startdate == d\}$; ▷ delivered drugs on date d
 $N := \{dr \mid \forall dr \in D \wedge dr.name \notin U.name\}$; ▷ newly delivered drugs
 $DC := \{dr \mid \forall dr \in D, \exists dr' \in U \text{ such that } dr.name ==$
 $dr'.name \wedge dr.dosage <> dr'.dosage\}$; ▷ dosage changed drugs
 $S := \{dr \mid \forall dr \in U \wedge dr.name \notin D.name \wedge dr.enddate < d\}$; ▷ recently
 stopped using drugs
 for each $d' \in U$ **do**
 if $\exists d'' \in D$ such that $d'.name == d''.name$ **then**
 $d' := d''$; ▷ update U with redelivered drugs
 $U := (U \setminus S) \cup N$; ▷ update U with newly delivered drugs
 $CD := N \cup DC \cup S$; ▷ considering drugs for calculating scores
 $CMD := CD.name \cap MDB$; ▷ considering main drugs
 $CSD := CD.name \cap SDB$; ▷ considering symptom-healing drugs
 $UD := CD.name \setminus (CMD \cup CSD)$; ▷ unclassified drugs
 $aScore = aScore + |CMD| \times w_{main} + |CSD| \times w_{symp} + |UD| \times w_{unk}$;
 Add $aScore$ to *scores*

3.4 Learning Group Treatment Regimens

The previous section has demonstrated our domain integrated algorithm which allows prescription recorded to be divided into periods based on the associated scores which reflect the change in the indication of prescribed drugs. In this section, we describe how a treatment regimen over a period of a given patient cluster is derived. We relax the chronological order of delivered drugs in a period and restrict the element of constructed treatment regimens to drug names only. Other information such as dosage, route, is assumed to be decided by the physicians.

The learned regimens were organized in a tree form. Starting from the root, we assign the most frequently prescribed drug d to its left child node and extract prescribed drugs excluding d of the patients who were treated by d . The drug assignment for next right child nodes will follow the similar approach applied on prescribed drugs of those patients who were not treated by left-hand side nodes in the same level. We recursively perform this procedure on internal nodes. To avoid learning too complicated details of the derived tree, we only perform the procedure until a certain level of the tree or when the number of patients treated by the most frequent drug for the parent node is still greater than a threshold. Algorithm 2 presents our ideas to construct the treatment regimen tree for a particular group of patients in a period.

Algorithm 2: Procedure for the construction of a treatment regimen tree

```

REGIMEN-Tree(depth, prescData, parent, traces)
  if prescData is empty or depth == maxDepth then
    ⊥ return
  d := most frequent drugs from prescData ;
  nPatients := number of patients who were treated by d;
  traces[parent, d] = “↖”;
  cNodePresc := prescribed records excluding d of patients treated by d ;
  rNodePresc := prescribed records of patients who were not treated by d ;
  if nPatients < threshold then
    | REGIMEN-Tree(depth, rNodePresc, parent, traces);
  else
    | REGIMEN-Tree(depth + 1, cNodePresc, d, traces);
    | REGIMEN-Tree(depth, rNodePresc, parent, traces);
  return ;

```

4 Experimental Evaluation

This section presents our experimental evaluation of the proposed framework for deriving typical treatment regimens from electronic medical records. The obtained results of the clustering analysis, treatment period identification and learned treatment regimen trees are also given and analyzed. Lastly, we propose a method to evaluate the efficacy of the derived treatment regimen trees in recommending prescribed drugs for new patients.

4.1 Experimental Design

Our experimental evaluation was performed on MIMIC III, a freely accessible critical care database [4]. We considered the treatment regimen of patients who were diagnosed with coronary artery disease as a case study. Since a patient could be diagnosed with multiple ICD codes, to ensure the homogeneity of our patient cohort, we only selected those whose primary ICD is coronary artery disease and comorbidity scores are zero for other disease groups. In addition, patients who were prescribed fewer than three times were also excluded from the experimental evaluation. The number of extracted patients is 707 of which we randomly selected 687 patients for training and left 20 patients for testing the efficacy of the learned treatment regimens. We followed the approach described in the Data Representation section to preprocess raw data. A summary of preprocessed data with illustration features is given in Table 1

We fit preprocessed data as input for MV.RBM with 200 hidden units since the trained error did not decrease significantly with a larger number of hidden units. The learned binary hidden states were then extracted as representation features for the subsequent clustering task. We employed hierarchical clustering with parameters are described in the previous section. For the task of treatment period identification, we extracted main drugs and symptom-healing drugs from

Kind of data (no.features)	Sample features (data type)
Demographic info (11)	age (numerical), gender (binary) admission type (categorical)
Laboratory data (175)	arterial blood pressure (numerical) atrial pacemaker(numerical)
Signs and symptoms data (1466)	abdominal discomfort (binary) ability to climb (binary) able to sleep (binary)

Table 1. A short summary of features in the dataset

DrugBank database. The referred typical symptoms of coronary artery disease in the literature are “heart attack”, “shortness of breath” and “chest pain”. We assigned the weight of main drugs, symptom healing drugs, and unclassified drugs to 1, 0.5 and 0.1, respectively. The threshold of cutting node in treatment learning algorithm was set to 10 patients. In our experiment, we derived regimens until a certain level of the tree. The depth parameter was set to 4.

4.2 Results

Fig. 3 describes a dendrogram of clustering results. It is noted that the trained patients themselves are homogeneous subjects in terms of diagnostic perspective. Thus, we preferred a relatively small number of clusters. Based on the visualization, we decided to group the trained patients into six clusters. The size of each cluster is 198, 69, 148, 43, 111 and 118 patients, respectively.

Fig. 4 presents a few randomly taken line charts of accumulative scores for eight patients. Interestingly, most of the plots follow similar patterns. There is a slight increase in scores at the beginning and the end of every treatment compared to the significant change at the center interval. Therefore, we decided to divide prescription records of trained patients into three periods. Fig. 5 illustrates an example of constructed regimen trees. Given a path in the tree, we note that the order of the nodes in this path should be understood as frequency order of drug use rather than chronological order of prescription time. It can be




Fig. 3. Dendrogram of hierarchical cluster analysis




Fig. 4. Sample line charts of accumulative scores of randomly taken patients. For each line chart, the vertical axis represents the accumulative scores, the horizontal axis represents the timestamps when a patient was prescribed.

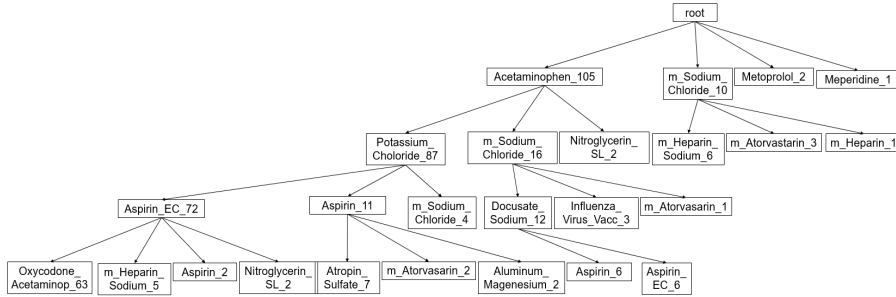


Fig. 5. Sample of learned regimen tree. The prefix “m” denotes for main drugs while the last integer denote the number of prescribed patients.

seen that the visualization can provide hint-drugs probably delivered together with a given drug. Therefore, physicians can use the learned trees as a checklist to decide which drugs are likely and unlikely to be prescribed.

4.3 Evaluation

We evaluate the efficacy of learned regimen trees in recommending prescribed drugs to new patients. It should be noted that patient records of the testing set are represented by the trained MV.RBM. We consider the patients in each resulting clusters as labeled data and assign the cluster index for test patients based on their nearest neighbors. Given a new patient p , let p' be his/her nearest neighbor which has been assigned to cluster c_i . The recommended drugs should be given to p in a particular period are drugs on the path of regimen tree of c_i in the same period such that p' was prescribed with each drug on that path.

Let T , nP , $\hat{D}_p^{t_j}$, and $D_p^{t_j}$ denote the test set, the number of periods, the recommended path, i.e., the set of recommended drugs for p over period t_j , and the set of prescribed drugs for p in that period, respectively. We propose two measures to evaluate the efficacy of learned regimens for the prescription recommendation task. These measures reflect how likely $\hat{D}_p^{t_j}$ is a subset of $D_p^{t_j}$. $\hat{D}_p^{t_j}$ is said “correct” if it is a non empty subset of $D_p^{t_j}$. In case $\hat{D}_p^{t_j}$ has non

empty intersection with $D_p^{t_j}$ but not its subset, we say the set $\hat{D}_p^{t_j}$ is “approximately correct”. We denote m_{cor} as the percentage of recommended paths which are “correct” and m_{app} as the percentage of recommended drugs actually prescribed in both “correct” and “approximately correct” paths. Let $I_A(B)$ define the indicator function which return 1 if $B \subset A$ or 0 otherwise. We have:

$$m_{cor} = \frac{1}{|T| \times nP} \sum_{p \in T} \sum_{j=1}^{nP} I_{D_p^{t_j}}(\hat{D}_p^{t_j}); \quad m_{app} = \frac{1}{|T| \times nP} \sum_{p \in T} \sum_{j=1}^{nP} \frac{|\hat{D}_p^{t_j} \cap D_p^{t_j}|}{|\hat{D}_p^{t_j}|}$$

We repeated our experiment 10 times for different training and testing sets. The obtained values of \bar{m}_{cor} and \bar{m}_{app} are 0.527 and 0.729, respectively. Although the obtained values of \bar{m}_{cor} should be further improved, to some extent, these measures show the efficacy of the regimen trees derived from our proposed framework.

5 Discussion

Comparing to related works in the literature, our work obtained more interesting results in terms of domain exploitation and knowledge representation. Rather than defining a similarity metric by a frequency-based approach for complex objects [9], we tracked the change of drug indication in prescribed drugs as a hint to discover treatment periods. It can be seen that the idea fits our natural thinking on detecting patients’ treatment periods given their prescription records. The common pattern found in Fig 4 has reconfirmed the rationality of our proposed domain-based algorithm. Moreover, representing the learned regimens in form of trees not only fully reflects the usage-frequency of drugs but also allows doctors to quickly recognize groups of frequently and infrequently prescribed drugs in each patient sub-cohort. Therefore, in terms of knowledge representation, it could be said that our work is superior to [2, 9] where the authors simply organized treatment patterns in flat form.

There are several reasons to explain the primitive results of our initial study on the task of treatment recommendation. Firstly, it is worth noting that we addressed the problem of treatment recommendation on MIMIC III, a practical and very challenging dataset. Even if it has been simplified to recommend in total up to 12 among many prescribed drugs for every patient, the problem is still not trivial as there are hundreds of different drugs given in the prescription records. Additionally, while our evaluation metrics directly assess whether the recommended drugs are prescribed to new patients, it is not clearly described in other studies how well the recommended treatments match the actual prescribed drugs. We leave the task of improving our prediction accuracy with a more deliberated framework for the future work.

6 Conclusion

In this paper, we have presented a generic framework to derive treatment regimens from electronic medical records. The proposed framework is novel in terms

of data utilization, domain incorporation, and regimen representation. The experimental evaluation has shown the efficacy of learned treatments for the task of prescription recommendation. Although further improvement should be made such as data cleaning and normalizing for clinical features, this study is a pioneering work which encourages researchers to exploit medical domain knowledge and address the treatment learning problem more thoroughly.

Acknowledgement

This work is partially sponsored by Asian Office of Aerospace R& D under agreement number FA2386-17-1-4094 and Vietnam National University at Ho Chi Minh City under the grant number B2015-42-02. We wish to thank Tu Dinh Nguyen for providing the implementation of the MV.RBM model.

References

1. Bodenreider, O.: The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research* **32**(suppl.1), D267–D270 (2004)
2. Huang, Z., Dong, W., Bath, P., Ji, L., Duan, H.: On mining latent treatment patterns from electronic medical records. *Data Mining and Knowledge Discovery* **29**(4), 914–949 (2015)
3. Huang, Z., Lu, X., Duan, H.: On mining clinical pathway patterns from medical behaviors. *Artificial intelligence in medicine* **56**(1), 35–50 (2012)
4. Johnson, A.E., Pollard, T.J., Shen, L., Lehman, L.w.H., Feng, M., Ghassemi, M., Moody, B., Szolovits, P., Celi, L.A., Mark, R.G.: MIMIC-III, a freely accessible critical care database. *Scientific data* **3** (2016)
5. Johnston, R.L., Carius, H.J., Skelly, A., Ferreira, A., Milnes, F., Mitchell, P.: A retrospective study of ranibizumab treatment regimens for neovascular age-related macular degeneration (namd) in Australia and the United Kingdom. *Advances in Therapy* **34**(3), 703–712 (2017)
6. Merhej, E., Schockaert, S., McKelvey, T.G., De Cock, M.: Recommending treatments for comorbid patients using word-based and phrase-based alignment methods. In: 29th Benelux Conference on Artificial Intelligence (BNAIC'17) (2017)
7. Savova, G.K., Masanz, J.J., Ogren, P.V., Zheng, J., Sohn, S., Kipper-Schuler, K.C., Chute, C.G.: Mayo clinical text analysis and knowledge extraction system (ctakes): architecture, component evaluation and applications. *Journal of the American Medical Informatics Association* **17**(5), 507–513 (2010)
8. Song, R., Wang, W., Zeng, D., Kosorok, M.R.: Penalized q-learning for dynamic treatment regimens. *Statistica Sinica* **25**(3), 901 (2015)
9. Sun, L., Liu, C., Guo, C., Xiong, H., Xie, Y.: Data-driven automatic treatment regimen development and recommendation. In: KDD. pp. 1865–1874 (2016)
10. Tamasauskas, D., Sakalauskas, V., Kriksciuniene, D.: Evaluation framework of hierarchical clustering methods for binary data. In: Hybrid Intelligent Systems (HIS), 2012 12th International Conference on. pp. 421–426. IEEE (2012)
11. Tran, T., Phung, D., Venkatesh, S.: Mixed-variate restricted Boltzmann machines. arXiv preprint arXiv:1408.1160 (2014)
12. Wishart, D.S., Knox, C., Guo, A.C., Cheng, D., Shrivastava, S., Tzur, D., Gautam, B., Hassanali, M.: DrugBank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic acids research* **36**(suppl.1), D901–D906 (2007)