Phenotypical Ontology Driven Framework for Multi-Task Learning

Mohamed Ghalwash
IBM Research
NY, USA
mohamed.ghalwash@ibm.com

Zijun Yao
IBM Research
NY, USA
zijun.yao@ibm.com

Prithwish Chakraborty
IBM Research
NY, USA
prithwish.chakraborty@ibm.com

James Codella
IBM Research
NY, USA
jvcodella@gmail.com

Daby Sow
IBM Research
NY, USA
sowdaby@us.ibm.com

ABSTRACT

Despite the large number of patients in Electronic Health Records (EHRs), the subset of usable data for modeling outcomes of specific phenotypes are often imbalanced and of modest size. This can be attributed to the uneven coverage of medical concepts in EHRs. We propose OMTL, an Ontology-driven Multi-Task Learning framework, that is designed to overcome such data limitations. The key contribution of our work is the effective use of knowledge from a predefined well-established medical relationship graph (ontology) to construct a novel deep learning network architecture that mirrors this ontology. This enables common representations to be shared across related phenotypes, and was found to improve the learning performance. The proposed OMTL naturally allows for multi-task learning of different phenotypes on distinct predictive tasks. These phenotypes are tied together by their semantic relationship according to the external medical ontology. Using the publicly available MIMIC-III database, we evaluate OMTL and demonstrate its efficacy on several real patient outcome predictions over state-of-the-art multi-task learning schemes. The results of evaluating the proposed approach on six experiments show improvement in the area under ROC curve by 9% and by 8% in the area under precision-recall curve.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning approaches:** **Multi-task learning.**

KEYWORDS

Embedding, Multitask, Knowledge-driven

ACM Reference Format:

Mohamed Ghalwash, Zijun Yao, Prithwish Chakraborty, James Codella, and Daby Sow. 2021. Phenotypical Ontology Driven Framework for Multi-Task Learning. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3450439.3451881>

1 INTRODUCTION

Patient data from large scale EHR are of increasing interest for clinical hypothesis generation and testing. These datasets can contain rich information regarding the medical history of millions of patients and cover varied aspects such as diagnosis, medication, and lab tests. However, these data can be noisy, sparse, and non-uniformly distributed such that studies on a specific outcome of interest can lead to very modest number of useful patient records. These data are often insufficient to learn the underlying complex disease phenomena, especially with deep learning models. The characteristics of data imbalance often result in a ‘vanishing data’ problem during the analysis of specific diseases of interest with low coverage. For instance, only 210 medical records (intensive care unit stays) can be used to study patients who have experienced Rheumatic disease of heart valve and Mitral valve stenosis on the well-known MIMIC-III database [17] that originally covers 42000 records present in the complete dataset. Analysis on such datasets can be impractical and suffer from significant lack of robustness. In another case, on an EHR database spanning more than 55 million patients, querying for patients within the age group 21 – 40 and having experienced concussions that led to mild cognitive impairments returns fewer than 1100 patients. Further characterization of this cohort to focus only on patients that eventually experienced dementia takes that number below 200.

Interestingly, a significant amount of domain knowledge relating medical concepts is readily available in several medical ontologies. Phenotypes [1] corresponding to observable physical properties of patients, including diseases, have been organized in such ontologies and expressed within EHR systems. We postulate that leveraging such phenotypical relations within EHR data to jointly learn multiple outcomes in a multi-task learning setting could provide solutions to this vanishing data problem and lead to more robust models.

In this paper, we introduce OMTL, an ontology driven framework for multi-task learning. The main tenant of our approach is to ground the structure of the proposed multi-task framework on phenotypical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM CHIL '21, April 8–10, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8359-2/21/04...\$15.00

<https://doi.org/10.1145/3450439.3451881>

medical concepts relations between patients from different cohorts. OMTL allows for the learning of outcomes for different phenotypes borrowing representation expressiveness from patients with similar phenotypes. This sharing allows the model to generate predictions for phenotypical cohorts even in phenotypes with small amounts of patient observations. Moreover, OMTL decouples the *outcome learning* from the *representation learning* of similar diseases. In this case, the model learns from patients with related diseases with different outcomes. For example, it can learn about mortality prediction for diabetic patients and readmission prediction for patients with hypertension by sharing information between diabetes and hypertension.

The contributions of this paper are 3 fold:

- (1) We propose an approach to model different phenotypical cohorts and outcomes using domain knowledge in the form of medical concept ontological graphs.
- (2) We propose a novel hierarchical deep learning architecture with gating mechanisms to specialize the outcome models of each phenotype using the ontological relationship.
- (3) We demonstrate the efficacy of the proposed OMTL method on the prediction of several outcomes on patients from MIMIC-III.

In the remainder of the paper, we start in Section 2 with a problem description, stating the limitations of existing state-of-the-art methods and motivating the needs for structural multi-task learning followed by describing the proposed OMTL method in Section 3. Next, in Section 4 we present the experimental results geared towards the evaluation of the gain achieved by leveraging the domain ontology to relate multiple tasks. Then we discuss specializing the OMTL architecture using reward shaping in Section 5. We follow with an overview of recent related works in Section 6, before summarizing our conclusions in Section 7 along with comments on the direction of future research.

2 PRELIMINARIES AND MOTIVATION

Let us take the following example as motivation for the proposed approach. Assume that we want to predict two outcomes, *in-hospital mortality* (we refer to as mortality hereafter) and *readmission*, for diabetic patients and also predict the *mortality* outcome for patients diagnosed with hypertension. A naive way to address this problem is to construct a cohort D^k for each individual outcome $k = \{1, 2, 3\}$; a cohort that includes patients with the diabetes phenotype who did (or did not) experience mortality; another cohort consists of diabetic patients with/without readmission, and lastly a cohort of patients diagnosed with hypertension and died or survived during the hospitalization. This is formulated as:

PROBLEM 1 (SINGLE-TASK LEARNING). *Given data D^k for task k , learn a predictive model for each task: $y = f^k(x) \forall x, y \in D^k$.*

Having limited data for each cohort reduces the opportunity to have an accurate model by learning each task individually. We thus propose to use multi-task framework to address these issues. Recently, a multi-task mixture of experts method was proposed to share information among tasks [21]. The model has a shared representation layer, called experts, to learn different mixture of representations. These representations are combined differently for

different tasks, which helps the model to share information among all tasks (outcomes) and at the same time to specialize the way these representations are aggregated for each task.

PROBLEM 2 (MULTI-TASK LEARNING). *Given data D^k for each task k , learn a predictive model for each task by sharing representations among tasks: $y = f^k(h(x)) \forall x, y \in D^k$ where $h(x)$ is a shared representation.*

A drawback of such approach is it does not assume the prediction of mortality and readmission for diabetic patients are more closely related to each other (because they are related to diabetic patients) than prediction of mortality in the hypertension cohort. Moreover, different cohorts of patients (diabetes, hypertension, etc.) might have different types of relationships. For example, it is known that diabetes and hypertension are closely interlinked because of similar risk factors, whereas diabetes and atrial fibrillation share common antecedents such as hypertension. Incorporating such semantic information could lead to a better representation specialized for each phenotype, which will subsequently help the prediction of outcomes of interest for those patients who are diagnosed with that phenotype.

On top of the aforementioned issues, there may not be enough data to effectively train such a multi-task deep learning model. As was discussed in the introduction, if the user is interested in a cohort of patients developing disorder of narrowing of the heart valves caused by rheumatic fever: “Rheumatic disease of heart valve” and/or “Mitral valve stenosis”, there are only 200 patients with these phenotypes in MIMIC-III database which are not enough to learn an accurate model, even after using multi-task framework. One possible solution is *data augmentation* by including cohorts of other diseases or phenotypes. In Section 6 we discuss several augmentation methods that have been previously proposed for this problem.

PROBLEM 3 (DATA AUGMENTATION). *Given data D^k for each task k , extract augmented data D' that share similar characteristics to the initial cohort.*

The results of applying the augmentation method is additional patients that can be included in the initial cohort. However, we are not particularly interested in predicting outcomes (mortality or readmission) for the additional cohort, we are only using these additional data to better train the predictive model. The challenges now become:

Challenge 1: how to include these additional patients who may not have the label of the outcome of interest?

Challenge 2: how to leverage the semantic relationship $\mathcal{G}$ among phenotypes of the augmented and the initial cohort?

PROBLEM 4 (ONTOLOGY-DRIVEN MULTI-TASK LEARNING). *Given data D^k for task k , augmented data D' and Ontology $\mathcal{G}$, learn a predictive model for each task: $y = f^k(h(D', \mathcal{G}, x)) \forall x, y \in D^k$.*

We focus to address these challenges in Problem 4. In particular, how to shape the function h such that it leverages the ontology and the augmented data to improve the accuracy performance of the predictive model. As evident by our six experiments, addressing these challenges boosts the AUC by up to 9% over state-of-the-art methods. In order to address these challenges, we propose to leverage

Figure 1: A directed acyclic subgraph $\mathcal{G}$ that encompasses nodes represented in the dataset. Each node represents one medical concept. The nodes in black (i.e., core nodes) are of specific interest to the user while the white nodes are added by the augmentation process in [12]. The number in each node is the SNOMED CT code for the specified phenotype.

a clinical ontology, such as SNOMED CT [10]. The ontology is represented as a directed acyclic graph $\mathcal{G}$ over of a set of nodes (e.g. diagnosis codes or phenotypes) and edges expressing parent-child *semantic* relationships. A sample of that ontology is shown in Figure 1 where black nodes represent phenotypes of the initial cohort and white nodes represent phenotypes of the augmented data. Including these semantic relationships would help the model to improve the phenotypical representation that can lead to more accurate predictions of outcomes specialized for each phenotype.

3 METHODS

The novelty of the proposed OMTL method is in two parts. First, we decouple the representation learning of each phenotype from the prediction of the outcome of interest as to allow sharing of representations for patients with the same phenotype for different predictive tasks. Second, we impose the ontology on the network architecture in order to share information among different but related phenotypes. These two parts are described in the following sections.

Let $X = \{x_i \in \mathbb{R}^d; \forall i\}$ be a matrix where x_i represents the feature vector for the record i with dimension d (e.g. vital signs and lab tests collected at the time of admission to the intensive care unit (ICU)). For each record i , we assume the presence of all phenotypical medical concepts tied to i , as it is commonly done in EHR systems. For example, all diagnoses the patient developed are known at the time of admission. Using that information, we construct a subgraph $\mathcal{G}$ of the ontology with only nodes corresponding to medical concepts that are present in X . Any outcome prediction can be formulated as a mapping from X into an outcome set Y . To enable sharing across outcomes, we break down this mapping as follows. We first define an encoding function $\mathcal{E} : X \rightarrow Z$, where Z is the representation space.

For this representation to be faithful, we enforce an unsupervised reconstruction process by learning a reconstruction function $\mathcal{R} : Z \rightarrow X$. The generated representation can be used as an input to a standard supervised machine learning algorithm f to predict Y , $f : Z \rightarrow Y$. Learning the parameters for the mappings $\mathcal{E}$, $\mathcal{R}$ and f can be achieved by minimizing $\mathcal{L}$:

$$\mathcal{L} = \underbrace{\sum_i \ell(f(\mathcal{E}(x_i)), y_i)}_{\mathcal{L}_1} + \lambda \underbrace{\sum_i \|\mathcal{R}(\mathcal{E}(x_i)) - x_i\|_2}_{\mathcal{L}_2} \quad (1)$$

where $\mathcal{L}_1$ is the total outcome prediction loss that aggregates the losses ℓ (e.g. cross entropy for classification and mean square error for regression) at the node levels corresponding to the potentially different outcomes. $\mathcal{L}_2$ is the reconstruction loss function. $\lambda \in \mathbb{R}^+$ is a hyper-parameter that weighs the importance of the two losses.

Learning a representation and a predictive model for one task is relatively straightforward by optimizing equation (1). However, learning such representations for multiple patient cohorts with different phenotypical profiles and multiple predictive tasks simultaneously is a challenging problem. We formalize OMTL by first describing a multi-task learner to decouple the outcome prediction and phenotypical representation in order to facilitate the incorporation of the hierarchical ontological structure as the second part of the model to address the aforementioned challenges.

3.1 Decoupling Outcome from Phenotypical Nodes

The OMTL model is depicted in Figure 2. Input data flows from the top into a mixture of E expert nodes responsible for clustering the input subspace to facilitate representation computations. We denote these expert nodes as $E_e(\cdot)$, $1 \leq e \leq E$, where each expert is a neural network producing a representation $E_e(x_i)$.

The outputs of these expert nodes are consumed by an array of P phenotypical representation nodes denoted $P_p(\cdot)$, $1 \leq p \leq P$. Each P_p maps directly into one medical concept (i.e., the number of representation nodes is equal to the number of nodes in the subgraph $\mathcal{G}$ defined in the previous section). Each of these P_p nodes ingest experts outputs via a gating function G_p to learn how to combine information from all expert nodes for the generation of representations for the current medical concept. The gating network G_p controls the combination of expert nodes to produce the input $\mathcal{M}_p$ to the representation node P_p (described in the green box in Figure 2). These networks are typically simple networks that produce combination weights for the experts at each representation node:

$$G_p(x_i) = \text{Softmax}(x_i \cdot W_p + b_p) \in \mathbb{R}^E$$

$$\mathcal{M}_p(x_i) = \sum_{e=1}^E G_p^e(x_i) E_e(x_i) \in \mathbb{R}^{d_e} \quad (2)$$

where $W_p \in \mathbb{R}^E$, $b_p \in \mathbb{R}$ are parameters to be learned by the gate network during training, and $G_p^e(x_i)$ is the entry e in the vector $G_p(x_i)$. The input $\mathcal{M}_p(x_i)$ is then passed to the representation node P_p to produce representations

$$\mathcal{E}_p(x_i) = P_p(\mathcal{M}_p(x_i)) \quad (3)$$

for the current medical concept p . The objective of this part is to dynamically (using the gates) combine multiple experts for the generation of useful representation for the node p .

The representations $\mathcal{E}_p(\cdot)$ are then passed as input to the *outcome* nodes. The outcome nodes leverage the representations to predict these outcomes. Each outcome node can be modeled as a neural network with the objective of predicting the outcome of interest. It

Figure 2: OMTL: the proposed hierarchical multi-task approach. The model decouples the outcome nodes (rectangle striped nodes) from phenotypical nodes (solid colored nodes). Input of the phenotypical nodes are controlled by two gating mechanism; one to weight the inputs from phenotypical parents (Gate H) and another one to control the input from the shared representation (Gate G). Each core node is associated with different outcome nodes. The output of the outcome node is the final prediction.

is quite important to note that *OMTL* can be optimized for multiple outcomes of interest as each representation node can have different and multiple outcomes, e.g. we can associate mortality prediction for certain medical concepts and readmission prediction for another (possibly overlapping) subset of medical concepts. In Figure 2, the phenotypical red node has two associated outcome nodes. This mimics the diabetic patients in our motivating example with two prediction tasks mortality and readmission.

The model is trained in a multi-task setting allowing expert nodes and representation nodes to share learned representations and support outcome nodes in their prediction tasks. The loss function to be optimized is

$$\mathcal{L} = \underbrace{\sum_i \sum_{\substack{p \in C \\ x_i \text{ has } p \text{ assoc. with } p}} \sum_{o \text{ with } p} \ell(f_p^o(\mathcal{E}_p(x_i)), y_i^o)}_{\mathcal{L}_1} + \lambda \underbrace{\sum_i \sum_{\substack{p \\ x_i \text{ has } p}} (\mathcal{R}_p(\mathcal{E}_p(x_i)) - x_i)^2}_{\mathcal{L}_2} \quad (4)$$

where $\mathcal{R}_p$ is the reconstruction network that reconstructs the input x_i from the representation $\mathcal{E}_p(x_i)$, ℓ is the prediction loss function and y_i^o is the label for the outcome o for the visit i . The second sum in $\mathcal{L}_2$ runs over all representation nodes (phenotypes) that x_i expresses forcing the representation from *all* nodes to be useful for the reconstruction of the original input. This allows the model to *incorporate augmented nodes* without labels and tackle **challenge 1** defined earlier. While the second sum in $\mathcal{L}_1$ runs over all *core* nodes C that x_i expresses to optimize the model for the predictive tasks for the initial cohort, the third sum runs over all outcome nodes o that are associated with the representation nodes p . Equation (4) allows

the model to utilize the augmented data (possibly unlabeled) and be optimized for the predictive tasks of the initial cohort (black nodes in Figure 1).

3.2 Imposing the Hierarchical Ontology

Optimizing equation (4) can be used to solve the multi-task problem outlined above. However, depending on the choice of the phenotypes of interest, there can be a significant variation in the number of patients spanned by the representation nodes. For example, nodes corresponding to a higher level concept (closer to the root) will typically span a much greater number of patients than the ones spanned by concepts at a leaf level of the ontology where concepts are very specialized. Interestingly, such nodes are semantically related via the medical concept ontology. Therefore, the generated representations across different medical concepts should share similar semantic relationships. *OMTL* aims to increase the information sharing among these nodes using the ontology.

As depicted in Figure 2, the representation nodes are now connected by mirroring the ontology subgraph $\mathcal{G}$. The representation nodes without ancestors (such as the green solid node) can be computed using the same gating principle outlined in equation (2). However, for a node with ancestors (such as the red node), the model shares both the representation from the expert nodes and its parents. Assuming that P_j has n parents, the resultant representation $\mathcal{E}_j$ at P_j is an aggregation of two vectors: (i) $\mathcal{M}_j$ computed using the combination of experts as given by equation (2) and, (ii) $\mathcal{P}_j$ computed using the combination of the representations from the parents as:

$$\begin{aligned} H_j(x_i) &= \text{Softmax}(x_i \cdot W_j + b_j) \\ \mathcal{P}_j(x_i) &= \sum_{k=1}^n H_j^k(x_i) \mathcal{E}_k(x_i) \end{aligned} \quad (5)$$

The input to the node P_j will be the aggregation of M_j and $\mathcal{P}_j$. The outcomes $\mathcal{E}_j(x_i)$ can be computed as:

$$\mathcal{E}_j(x_i) = P_j(M_j(x_i), \mathcal{P}_j(x_i)) \quad (6)$$

Using this strategy, the ancestors of a representation node will be less affected by the ‘less data’ problem and the structure thereby acts as a forcing mechanism to help the network to learn the representation nodes and outcome with higher efficacy. This allows OMTL to tackle **challenge 2**. During the forward and backward pass of the training, data x_i is only routed to representation nodes corresponding to phenotypes expressed in x_i . Algorithm 1 summarizes how OMTL process generates representations and predictions according to the proposed learning architecture.

Algorithm 1: OMTL

Input : data record x_i

```

1 for each node  $j$  do {nodes are sorted based on graph level}
2   if  $x_i$  expresses medical concept  $j$  then
3      $M_j \leftarrow$  apply equation (2)
4      $\mathcal{E}_j \leftarrow$  apply equation (3)
5   if  $j$  has parents then
6      $\mathcal{P}_j \leftarrow$  apply equation (5)
7      $\mathcal{E}_j \leftarrow$  apply equation (6)
8   for each outcome  $o$  associated with node  $j$  do
9      $\hat{y}^o = f_j^o(\mathcal{E}_j)$ 

```

4 RESULTS

In this section, we present an experimental evaluation of OMTL focusing on the impact provided by the ontology inspired architecture of OMTL that we learn within a multi-task framework. We have performed these experiments on cohorts of real patient data described in the next section. We then present details on the experimental setup before presenting the actual experimental results applying OMTL on the constructed cohorts.

4.1 Data Description and Cohort Construction

Data were extracted from the publicly available Medical Information Mart for Intensive Care (MIMIC-III) database [17], which is a large single-center database of patients admitted to critical care units at a large tertiary hospital. While MIMIC-III includes several facets of patient medical records, we focus our experiments on the vital signs, observations and diagnosis codes following the proposed benchmark cohort construction described in [14]. Since MIMIC-III tracks diagnosis codes in ICD-9 format, only preserving a shallow two-level hierarchical structure of medical concepts, we mapped all the observed ICD-9 diagnosis codes to SNOMED CT codes to construct a graph with a sufficiently deep hierarchy of medical concepts¹. As a result, we obtained a complex knowledge graph where each node is associated with a set of ICU stays.

We also followed the data augmentation process described in [12] to obtain three ontology graphs: Liver Disorder (LD), Heart

Disorder (HD), and Kidney Disorder (KD). For each graph, two outcomes were considered; prediction of in-hospital mortality and *Disorders of lipid metabolism* (DLM) for the LD graph, *Acute myocardial infarction* (AMI) for the HD graph, and *Chronic kidney disease* (CKD) for the KD graph. More details can be found in the supplementary. We followed [14] for processing the data and outcomes to predict. The corresponding data are multivariate time series for each ICU stay. Following [14], we impute missing data and collapse time series data to a single time point using the mean values for each ICU stay. Statistical details on the selected cohorts can be found in Table 1.

Table 1: Statistics of datasets used for experiments. Graph name (# of nodes / # core nodes). Number of visits (number of positive labels).

Graph	Task	# Visits (+)	Prevalence
LD (26/5)	Mortality	4767 (837)	0.175
	DLM	9073 (1474)	0.162
HD (21/3)	Mortality	13067 (1867)	0.143
	AMI	24135 (4039)	0.167
KD (78/1)	Mortality	20316 (2749)	0.135
	CKD	39108 (5588)	0.143

4.2 Experimental Setup

Our primary objective is to test the hypothesis that adding the structure information from the ontology helps the model to produce more accurate predictions on the phenotypes of interest. To test this hypothesis, we propose to compare OMTL against multi-task baselines. Multi-task mixture of expert models [21] is the state-of-the-art multi-task model, which uses mixture of experts [16] to share information among *all* outcomes without using any prior knowledge about the semantic phenotypical relationship. Moreover, the baseline model is optimized for supervised learning, i.e. it assumes that labels for all patients are provided, therefore the model will not benefit from the unlabeled augmented data. To have a fair comparison, we modified the loss function of the baseline model by using the loss structure outlined in equation (4). This allows the model to benefit from the augmented data. This model is referred to as MMOE in the rest of this section. In addition, we compare the proposed method to mixture of experts (MOE) [11] where multiple experts were utilized but no gates, and the most commonly used shared-bottom (SB) multi-task deep learning structure [3] where there is only one expert.

OMTL utilizes a gating mechanism H enforcing the ontological structural relationship across representation nodes. To train OMTL, we first train it without considering the hierarchy to learn the experts and their corresponding gates G . Then, we freeze these components and subsequently fine-tune for the outcomes and representations at the nodes by admitting the hierarchical structure.

Implementation details: Each module in the deep learning structure (expert, representation, reconstruction, or outcome node) is implemented as a one layer feed-forward neural network. We used 3 experts for MOE, MMOE, and OMTL. The size of the expert, reconstruction, representation modules are 41×5 , 5×41 ,

¹https://www.nlm.nih.gov/research/umls/mapping_projects/icd9cm_to_snomedct.html

5×5 , respectively. The activation functions are Leaky ReLU for experts, Softplus for the representation layers, and ReLU for the reconstruction. Each expert has a dropout layer for regularization. The size of the gate G_p is 41×3 and the size of the gate H_j is $41 \times n_j$ where n_j is the number of parents for the representation node j . We performed hyper-parameter tuning and found good performance and stability across all models by setting the batch size to 64, learning rate to 0.001, dropout to 0.5, λ to 0.0001, and using Adam as the optimizer [18]. We evaluated the models using stratified 5-fold cross validation. the cross validation was done so as to ensure we have the same distribution of positive and negative examples in each node. All models were evaluated on exactly the same folds for a fair comparison.

4.3 Multi-task Outcome Results

We compare all baselines with OMTL using area under ROC curve (AUC-ROC). The results are shown in Table 2. We also compare all models using average precision recall (APS) which summarizes the precision-recall curve; the results are shown in Table 3. To test the significance of the OMTL model, we applied the DeLong test to test whether the ROC curves are statistically significant [9, 28]. Whenever OMTL is statistically significant (pvalue < .05) than the baseline, we superscript the baseline with *.

Table 2: AUC under ROC for all models. * indicates when the OMTL model is statistically significant than the corresponding baseline.

Graph	Task	SB	MOE	MMOE	OMTL
LD	Mortality	0.715	0.742*	0.715*	0.750
	DLM	0.572*	0.592*	0.576*	0.626
HD	Mortality	0.723	0.700	0.713	0.729
	AMI	0.535*	0.537*	0.561*	0.593
KD	Mortality	0.728	0.674	0.697	0.728
	CKD	0.743	0.736	0.749	0.749

Table 3: APS using PRC for all models.

Graph	Task	SB	MOE	MMOE	OMTL
LD	Mortality	0.414	0.448	0.429	0.478
	DLM	0.086	0.106	0.084	0.109
HD	Mortality	0.211	0.221	0.242	0.262
	AMI	0.180	0.182	0.202	0.214
KD	Mortality	0.252	0.219	0.233	0.267
	CKD	0.806	0.798	0.806	0.806

Analysis: We can see from Table 2 that the gain in AUC-ROC that the structure is adding to MMOE ranges from 5%-9% for the liver disorder data with ontological graph shown in Figure 3, and the gain in APS ranges from 11%-30% (Table 3). It is worth noting that the core nodes in this graph are well-connected to each other. We hypothesize that the benefit of using the structural information

among such well-connected core nodes helps OMTL to attain a significant boost in sharing information and hence provides more accurate predictions than MMOE.

Figure 3: Ontological graphs populated with ICU stays from MIMIC-III for the liver diseases. Core nodes are shown in red, and augmented nodes are shown in blue.

The AUC-ROC gain in the heart disorder graph (Figure 4) ranges from 2% to 6%, and the APS gain ranges from 6%-8%. The gain is not as significant as in the liver disorder case. A possible explanation for this is that the core nodes are less connected to each other. In such situation, the training may benefit by leveraging more outcome nodes. Similar conclusions can be reached in the kidney graph (Figure 5) where there is only one core node and the AUC-ROC (and APS) for OMTL is similar to the MMOE model. In the next section, we discuss more this hypothesis and show how to further improve OMTL by tweaking the loss function.

Figure 4: Ontological graphs populated with ICU stays from MIMIC-III for the heart disorder. Core nodes are shown in red, and augmented nodes are shown in blue.

Finally, since all datasets are highly imbalanced (Table 1), the APS results in Table 3 indicates that OMTL is able to predict the rare positive class more accurately than MMOE, which is very important in medical domain because failing to identify patients at risk is very serious since a disease has been ignored.

5 DISCUSSION: SPECIALIZING THE OMTL ARCHITECTURE

As discussed in the experimental setup, we trained the models by reducing the reconstruction loss at all nodes and optimizing for the outcome loss at only the core nodes of interest C . This method of training ensures that OMTL can be used in scenarios

Figure 5: Ontological graphs populated with ICU stays from MIMIC-III for the kidney disorder. Core nodes are shown in red, and augmented nodes are shown in blue.

where outcomes of interest at individual representation nodes can be different. In such a training scenario, the information content at the augmented nodes is only used for learning the experts by ensuring that the model is able to reconstruct a broader variety of phenotypical information. However, when all the nodes of interest are concerned with *the same outcome*, we can fine tune the entire architecture to be discriminative w.r.t. that outcome (i.e. by reducing the outcome loss at the augmented nodes during training).

We devised an additional set of experiments to test the impact of optimizing OMTL on the core and augmented nodes together. For this set of experiments, we can control the importance of the augmented nodes towards the discriminative power of the network by weighing the losses from the nodes differently. The loss for the network for the outcome o is given as:

$$\text{loss}(\mathcal{G}, o) \propto \sum_{k \in \mathcal{G}} w_k \ell(y_k^o, \hat{y}_k^o) \quad (7)$$

where, y_k^o is the outcome o for all records at the node k , $\hat{y}_k^o$ is the prediction, and w_k is a weight to modulate the losses that control the reward the network gets for discriminating at the node k . There could be many ways of shaping such rewards based on a reward function $\mathcal{W}$. Intuition and domain knowledge can be used to define the right structure for $\mathcal{W}$. In this paper, we explore a linear reward scheme controlled by the level l_k of the node k as given below:

$$w_k = \mathcal{W}(k) \propto \left(\frac{l_k}{d}\right)^f, \quad f \in (-1, 1) \quad (8)$$

where d is the depth of $\mathcal{G}$ and f is a scaling factor. When $f = 1$, we reward the network for being more accurate at the leaves of $\mathcal{G}$ in contrast to rewarding the network for being accurate at the root levels when $f = -1$. $f = 0$ implies equal rewards at all nodes. After optimizing OMTL accordingly for different values of f , we evaluate its performance on the core nodes. We present the results of fine-tuning OMTL with reward shaping as expressed in equation (8) on Table 4.

As can be seen, including the outcome loss for the augmented nodes helps in specializing the OMTL model. Additionally, different forms of reward shaping can affect the performance of OMTL. For example, for the heart disorder mortality, the AUC of OMTL has been improved from 0.729 (see Table 2) when optimizing the model on only the core nodes to 0.782 when optimizing the model for both core and augmented nodes. Improvements are also observed for the other experiments. This verifies our hypothesis that the connectivity

Table 4: Specializing the OMTL model at different levels in the hierarchy. $f = 1$ rewards the network for being more accurate at the leaves. $f = -1$ rewards the network for being more accurate at the root levels. $f = 0$ implies equal rewards at all levels.

Group	Task	$f = 1$	$f = 0$	$f = -1$
LD	Mortality	0.771	0.783	0.781
	DLM	0.663	0.652	0.656
HD	Mortality	0.778	0.782	0.781
	AMI	0.622	0.612	0.625
KD	Mortality	0.792	0.793	0.795
	CKD	0.766	0.765	0.765

of the graph assists the model in specializing for the outcome of interest.

6 RELATED WORK

Multi-task learning: Multi-task learning has received a significant amount of attention in recent years. Notably, [21] adapts a mixture-of-experts structure to jointly learn multiple outcomes by sharing a single expert, and leverages gating networks to specialize for each outcome and joint representation learning. In [2, 15], the authors employ a two-level hierarchy to regularize a learned model to share information among multiple related outcomes. [35] extended these approaches with multi-level hierarchies using a cascade approach to progressively learn outcomes at prior levels to improve learning on outcomes in subsequent levels. The model proposed in [25] learns intermediate representations for predictive tasks at multiple scales for the same input via a cascade of connected layers. In [26], the authors proposed a way to incorporate a set of *carefully* selected semantic NLP tasks into the multi-task model. “Easy” to learn tasks are supervised at the lower levels of the architecture while more complex interactions are kept at deeper layers. A soft-ordering of the shared layers was proposed in [22] to learn both layers and their order on each task.

In most cases, the shared layers are integrated in parallel order, where the outcome of one layer is fed into the next layer. OMTL is motivated by these efforts but learns phenotype specific predictive models in a multi-task setting to share information across outcomes. Most importantly, OMTL decouples the outcome nodes from the phenotypical representation nodes. Representations learned for different phenotypes can be further specialized w.r.t. different tasks (e.g. readmission and mortality prediction) by leveraging the ontology. In addition, OMTL has the flexibility to be optimized for only a set of phenotypes in the graph (core nodes) and assign different weights to different nodes based on the domain-knowledge hierarchy.

Knowledge graph: Structured domain knowledge, often in graph-form, has been utilized to achieve more robust inferences over purely data-driven models. For example, [31] present a framework to enable the use of various kinds of external knowledge bases to retrieve relevant answers to a given question. It provides a graph-based model which maps text phrases to concepts in the knowledge graph

using several strategies. Finally, the extracted graph is then used as an input to a neural network model for entailment.

Outcome prediction from EHR:

Over the years, researchers have developed models on EHR data for the prediction of adverse events such as early identification of heart failure [29], readmission prediction [32], and acute kidney injury prediction [30]. Other notable works include [33] where the authors presented a multi-level attention mechanism to derive patient specific representations. The authors split the trajectory of patients into sub-sequences and use within-subsequence and between-subsequence attention levels sequentially to generate robust patient representations. Furthermore, in [13] the authors present four clinical prediction benchmarks using data derived from the MIMIC-III database, including In-hospital mortality, Decompensation, Length of stay, and Acute care phenotype classification. Additionally, they propose a deep learning model called multi-task RNN to empirically validate the four prediction benchmarking tasks. This benchmark is used as the basis for evaluating OMTL.

Recently, there has been a great interest in representing longitudinal patient data from EHR systems as low-dimensional vectors or embeddings to capture various aspects of the patient's health. One of the primary works in this field is [5] where the authors proposed a multi-layer representation learning tool for learning code and visit representations from an EHR dataset. [6] investigated the problem of data scarcity in healthcare and proposed to supplement EHR datasets with a medical ontology to provide semantic structure. The proposed attention based approach called GRAM, learns the representation of a medical concept by combining information from the concept's ancestors in the ontology. Subsequently, [8] proposed a method called MiME, to use the medical ontology to generate more robust patient embeddings by imposing a multi-level hierarchy on EHR data defined by an ontology. The levels include visits, diagnosis, and treatments.

While our proposed OMTL model also tackles the data scarcity problem by leveraging medical ontologies, it differs from these works in how the ontologies are used in learning representations. In GRAM, patient representations are learned as weighted combination of the medical concepts in the ontology. Similarly in MiME, robust patient embeddings are attained by considering the hierarchy of medical events. However, OMTL is aimed at learning the patient representations for multiple phenotypes in a joint multi-task manner i.e. the same patient can have different representations depending on the phenotype of interest but still benefits from the information present in more common phenotypes.

Data Augmentation: Several data augmentation algorithms were proposed but mainly outside of the healthcare domain [19, 24, 27]. Deep learning approaches have been proposed to generate synthetic EHR data [4, 7, 20]. These models might be a source of additional biases and therefore are less desirable. Patient similarity [23, 34] metrics are also used for data augmentation but often lack to take into account the semantic relationship between diseases the patients have. ODVICE was proposed in [12] that uses Monte-Carlo algorithm for data augmentation from real data. The augmented data have semantic relationship to the initial cohort. This is the algorithm that was used in our experiments.

7 CONCLUSIONS

We proposed OMTL, a novel multi-task learning approach to analyze EHR data. OMTL models outcomes by incorporating knowledge from an existing ontology of phenotypical information in a multi-task setting. We have demonstrated the utility of this approach by comparing it against state-of-the-art knowledge agnostic multi-task framework over a set of experiments conducted on real patients data. Our ongoing efforts are aimed at deploying the framework for large scale EHR analysis and exploring its ability to generate reusable robust patient representations at phenotype levels. Also, the proposed method is agnostic to the EHR application, so in the future we seek to apply the method beyond EHRs in any other application where domain knowledge (or ontology) is available.

REFERENCES

- [1] Definition phenotype. <https://www.nature.com/scitable/definition/phenotype-phenotypes-35>. Accessed: 2019-05-22.
- [2] Amr Ahmed, Abhimanyu Das, and Alexander J Smola. 2014. Scalable hierarchical multitask learning algorithms for conversion optimization in display advertising. In *Proceedings of the 7th ACM international conference on web search and data mining*. ACM, 153–162.
- [3] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [4] Zhengping Che, Yu Cheng, Shuangfei Zhai, Zhaonan Sun, and Yan Liu. 2017. Boosting deep learning risk prediction with generative adversarial networks for electronic health records. In *2017 IEEE International Conference on Data Mining (ICDM)*. IEEE, 787–792.
- [5] Edward Choi, Mohammad Taha Bahadori, Elizabeth Searles, Catherine Coffey, Michael Thompson, James Bost, Javier Tejedor-Sojo, and Jimeng Sun. 2016. Multi-layer Representation Learning for Medical Concepts. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1495–1504.
- [6] Edward Choi, Mohammad Taha Bahadori, Le Song, Walter F Stewart, and Jimeng Sun. 2017. GRAM: graph-based attention model for healthcare representation learning. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 787–795.
- [7] Edward Choi, Siddharth Biswal, Bradley Malin, Jon Duke, Walter F Stewart, and Jimeng Sun. 2017. Generating multi-label discrete patient records using generative adversarial networks. *arXiv preprint arXiv:1703.06490* (2017).
- [8] Edward Choi, Cao Xiao, Walter Stewart, and Jimeng Sun. 2018. MiME: Multilevel Medical Embedding of Electronic Health Records for Predictive Healthcare. In *Advances in Neural Information Processing Systems*. 4552–4562.
- [9] Elizabeth R DeLong, David M DeLong, and Daniel L Clarke-Pearson. 1988. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics* (1988), 837–845.
- [10] Kevin Donnelly. 2006. SNOMED-CT: The advanced terminology and coding system for e-Health. *Studies in health technology and informatics* 121 (2006), 279.
- [11] David Eigen, Marc'Aurelio Ranzato, and Ilya Sutskever. 2013. Learning factored representations in a deep mixture of experts. *arXiv preprint arXiv:1312.4314* (2013).
- [12] Mohamed Ghalwash, Zijun Yao, Prithwish Chakraborty, James Codella, and Daby Sow. 2020. ODVICE: An Ontology-Driven Visual Analytic Tool for Interactive Cohort Extraction. *arXiv preprint arXiv:2005.06434* (2020).
- [13] Hrayr Harutyunyan, Hrant Khachatryan, David C Kale, and Aram Galstyan. 2017. Multitask learning and benchmarking with clinical time series data. In *MLAH: Machine Learning for Health workshop, held at Advances in Neural Information Processing Systems*.
- [14] Hrayr Harutyunyan, Hrant Khachatryan, David C. Kale, Greg Ver Steeg, and Aram Galstyan. 2019. Multitask learning and benchmarking with clinical time series data. *Scientific Data* 6, 1 (2019), 96. <https://doi.org/10.1038/s41597-019-0103-9>
- [15] Jingrui He and Yada Zhu. 2012. Hierarchical multi-task learning with application to wafer quality prediction. In *2012 IEEE 12th International Conference on Data Mining*. IEEE, 290–298.
- [16] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, Geoffrey E Hinton, et al. 1991. Adaptive mixtures of local experts. *Neural computation* 3, 1 (1991), 79–87.
- [17] 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 (2016), 160035.
- [18] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [20] Fenglong Ma, Yaqing Wang, Jing Gao, Houping Xiao, and Jing Zhou. 2020. Rare Disease Prediction by Generating Quality-Assured Electronic Health Records. In *Proceedings of the 2020 SIAM International Conference on Data Mining*. SIAM, 514–522.
- [21] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. 2018. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, 1930–1939.
- [22] Elliot Meyerson and Risto Miikkulainen. 2018. Beyond shared hierarchies: Deep multitask learning through soft layer ordering. In *6th International Conference on Learning Representations (ICLR)*.
- [23] Enea Parimbelli, Simone Marini, Lucia Sacchi, and Riccardo Bellazzi. 2018. Patient similarity for precision medicine: A systematic review. *Journal of biomedical informatics* 83 (2018), 87–96.
- [24] Holger R Roth, Christopher T Lee, Hoo-Chang Shin, Ari Seff, Lauren Kim, Jianhua Yao, Le Lu, and Ronald M Summers. 2015. Anatomy-specific classification of medical images using deep convolutional nets. In *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 101–104.
- [25] Ramon Sanabria and Florian Metze. 2018. Hierarchical Multi Task Learning With CTC. In *Workshop On Spoken Language Technology, held at IEEE SLT*. IEEE.
- [26] Victor Sanh, Thomas Wolf, and Sebastian Ruder. 2019. A hierarchical multi-task approach for learning embeddings from semantic tasks. In *AAAI*.
- [27] Patrice Y Simard, Dave Steinkraus, and John C Platt. 2003. Best practices for convolutional neural networks applied to visual document analysis. In *null*. IEEE, 958.
- [28] Xu Sun and Weichao Xu. 2014. Fast implementation of DeLong’s algorithm for comparing the areas under correlated receiver operating characteristic curves. *IEEE Signal Processing Letters* 21, 11 (2014), 1389–1393.
- [29] Jimeng Sun et. al. 2012. Combining knowledge and data driven insights for identifying risk factors using electronic health records. In *AMIA Annual Symposium Proceedings*, Vol. 2012. American Medical Informatics Association, 901.
- [30] Nenad Tomašev, Xavier Glorot, Jack W Rae, Michal Zielinski, Harry Askham, Andre Saraiva, Anne Mottram, Clemens Meyer, Suman Ravuri, Ivan Protsyuk, et al. 2019. A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature* 572, 7767 (2019), 116.
- [31] Xiaoyan Wang, Pavan Kapanipathi, Ryan Musa, Mo Yu, Kartik Talamadupula, Ibrahim Abdelaziz, Maria Chang, Achille Fokoue, Bassem Makni, Nicholas Mattei, et al. 2018. Improving Natural Language Inference Using External Knowledge in the Science Questions Domain. *arXiv preprint arXiv:1809.05724* (2018).
- [32] Cao Xiao, Tengfei Ma, Adji B Dieng, David M Blei, and Fei Wang. 2018. Readmission prediction via deep contextual embedding of clinical concepts. *PLoS one* 13, 4 (2018), e0195024.
- [33] Jinghe Zhang, Kamran Kowsari, James H Harrison, Jennifer M Lobo, and Laura E Barnes. 2018. Patient2Vec: A Personalized Interpretable Deep Representation of the Longitudinal Electronic Health Record. *IEEE Access* 6 (2018), 65333–65346. <https://doi.org/10.1109/ACCESS.2018.2875677>
- [34] Zihao Zhu, Changchang Yin, Buyue Qian, Yu Cheng, Jishang Wei, and Fei Wang. 2016. Measuring patient similarities via a deep architecture with medical concept embedding. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 749–758.
- [35] Alon Zweig and Daphna Weinshall. 2013. Hierarchical Multi-Task Learning: a Cascade Approach Based on the Notion of Task Relatedness. In *Theoretically Grounded Transfer Learning workshop, held at International Conference on Machine Learning (ICML)*.

SUPPLEMENTARY

Cohort construction: The ODVICE method [12] was proposed to expand the cohort of interest by leveraging the SNOMED CT clinical ontology [10]. Starting from the *core* phenotypes the user is interested in, the method leverages the ontology to guide the selection of additional patients records from semantically-related nodes $\mathcal{G}$. The method starts with a small subgraph that constitutes the core nodes. Then, it employs a Monte-Carlo algorithm to expand the graph by adding nodes that are within a proximity (to ensure semantic relationship to the initial cohort) to the initial subgraph. The additional nodes are selected based on two criteria: (1) its proximity to the initial subgraph, and (2) how similar the characteristic of

the additional patients to the initial cohort. For example, including patients diagnosed with “Heart disease” would help the model to learn a good representation for patients diagnosed with “Rheumatic disease of heart valve” as both of these phenotypes are semantically related based on SNOMED CT ontology.

Given the knowledge graph we generate three ontology graphs: Heart Disease (HD), Liver Disease (LD) and Kidney Disease (KD). For each graph, we start with a few core nodes (e.g., specific SNOMED nodes identified by domain experts for a certain phenotype) and include more SNOMED nodes to augment the ICU stays for analysis by applying the ODVICE method [12]. The data augmentation algorithm starts from this set of core nodes and grows the set by choosing a series of predecessors using an MCMC sampler with a maximum of 2 hops away from the core nodes. We apply this node growth procedure for a predefined number of iterations to obtain connected graphs with sufficient nodes and stays for training.

The LD Graph: One of the prediction tasks for the core nodes in the LD graph is to predict whether the patient developed “disorder of lipid metabolism” (this is not of the nodes in the given graph). We started from 5 core nodes (highlighted in Table 6). Those 5 nodes contain 1115 records which are not enough to learn a deep learning model. We expand those nodes by applying the method in [12] to obtain a final graph of 26 nodes that has 9081 records. The final LD graph is shown in Figure 3. Those additional nodes reveals related knowledge to improve the tasks on core nodes such as “Drug-related disorder”, “Poisoning”, and “Disorder of liver”. Adding additional patients who diagnosed with these phenotypes might help to learn more informed representation for patients diagnosed with the core nodes and be subsequently beneficial for predicting “disorder of lipid metabolism”. As clear from the graph, the nodes are very well connected which helps OMTL to significantly boost the AUC for the predictive tasks for the core nodes.

The HD Graph: The heart disease graph (Figure 4) has 21 nodes, three of them are the core nodes which has 3210 unique stays. The nodes are listed in Table 7

The KD Graph: The kidney disease graph (Figure 5) has entirely 39151 stays spanning 78 nodes, with only 1007 stays in the core node. The nodes are listed in Table 5.

Table 5: Medical concepts of kidney disease graph. There are 1007 and 39151 unique stays of core nodes and entire graph respectively. Highlighted rows are the core nodes.

No.	SNOMED	Description	Stays	No.	SNOMED	Description	Stays
0	116223007	Complication (disorder)	6555	39	42399005	Renal failure syndrome (disorder)	10419
1	118940003	Disease of nervous system (disorder)	13226	40	424736006	Diabetic peripheral neuropathy (disorder)	1279
2	127013003	Diabetic renal disease (disorder)	1007	41	429490004	Disorder affecting transplanted structure (dis...	632
3	127321000	Acute digestive system disorder (disorder)	1587	42	43064006	Hydronephrosis (disorder)	323
4	14669001	Acute renal failure syndrome (disorder)	7550	43	431856006	Chronic kidney disease stage 2 (disorder)	104
5	16934004	Renal vascular disorder (disorder)	286	44	431857002	Chronic kidney disease stage 4 (disorder)	284
6	193003	Benign hypertensive renal disease (disorder)	136	45	433144002	Chronic kidney disease stage 3 (disorder)	609
7	193186008	Polyneuropathy in disease EC (disorder)	1280	46	44054006	Diabetes mellitus type 2 (disorder)	8475
8	19660004	Disorder of soft tissue (disorder)	21819	47	441574008	Atherosclerosis of artery (disorder)	2080
9	197679002	Glomerular disease (disorder)	198	48	441656006	Hyperglycemic crisis in diabetes mellitus (dis...	218
10	230572002	Diabetic neuropathy (disorder)	1279	49	443820000	Disorder of kidney and/or ureter (disorder)	13504
11	234348004	Anemia of renal disease (disorder)	1134	50	444691002	Necrosis of cortex of kidney (disorder)	1661
12	236011003	Retroperitoneal cyst (disorder)	167	51	444794000	Acute necrosis of cortex of kidney (disorder)	1621
13	236423003	Renal impairment (disorder)	10859	52	449759005	Complication of systemic hypertensive disorder...	3212
14	236424009	Acute renal impairment (disorder)	7426	53	45281005	Atherosclerosis of renal artery (disorder)	191
15	236437007	Structural and functional abnormalities of the...	165	54	46177005	End stage renal disease (disorder)	1504
16	236439005	Cystic disease of kidney (disorder)	165	55	46635009	Diabetes mellitus type 1 (disorder)	591
17	237620003	Abnormal metabolic state in diabetes mellitus ...	224	56	473449006	Disease of non-coronary systemic artery (disor...	3222
18	237621004	Diabetic severe hyperglycemia (disorder)	218	57	4855003	Diabetic retinopathy (disorder)	734
19	238981002	Soft tissue complication of diabetes mellitus ...	1279	58	49455004	Diabetic polyneuropathy (disorder)	1279
20	25093002	Diabetic oculopathy (disorder)	759	59	49708008	Anemia of chronic renal failure (disorder)	1134
21	271737000	Anemia (disorder)	11288	60	51292008	Hepatorenal syndrome (disorder)	219
22	29555009	Retinal disorder (disorder)	985	61	52845002	Nephritis (disorder)	154
23	310647000	Anemia secondary to renal failure (disorder)	1134	62	56051008	Ketoacidosis (disorder)	218
24	35455006	Acute tubular necrosis (disorder)	1621	63	58574008	Acute nephropathy (disorder)	8957
25	36171008	Glomerulonephritis (disorder)	153	64	58797008	Complication of transplanted kidney (disorder)	403
26	363169009	Inflammation of specific body organs (disorder)	15032	65	609410002	Necrosis of anatomical site (disorder)	5689
27	367991000119101	Hyperglycemia due to type 1 diabetes mellitus ...	178	66	709044004	Chronic kidney disease (disorder)	4590
28	373406006	Inflammatory disorder of genitourinary system ...	590	67	721283000	Acidosis due to type 1 diabetes mellitus (diso...	178
29	38481006	Hypertensive renal disease (disorder)	3079	68	722223000	Cyst of kidney (disorder)	167
30	390834004	Non proliferative diabetic retinopathy (disorder)	714	69	73211009	Diabetes mellitus (disorder)	10195
31	420270002	Ketoacidosis in type I diabetes mellitus (diso...	178	70	74627003	Diabetic complication (disorder)	2583
32	420422005	Ketoacidosis in diabetes mellitus (disorder)	218	71	75934005	Metabolic disease (disorder)	22939
33	420868002	Disorder associated with type I diabetes mell...	464	72	767094002	Disorder of renal parenchyma (disorder)	1830
34	421326000	Neurologic disorder associated with type II di...	913	73	86249007	Obstructive nephropathy (disorder)	324
35	421468001	Neurological disorder associated with type I d...	336	74	90688005	Chronic renal failure syndrome (disorder)	3987
36	422014003	Disorder associated with type II diabetes mell...	1074	75	90708001	Kidney disease (disorder)	12881
37	422088007	Neurologic disorder associated with diabetes m...	1522	76	95568003	Renal tubular disorder (disorder)	1644
38	422099009	Diabetic oculopathy associated with type II di...	354	77	95570007	Kidney stone (disorder)	172

Table 6: Medical concepts of liver disease graph. There are 1115 and 9081 unique stays of core nodes and entire graph respectively. Highlighted rows are the core nodes.

No.	SNOMED	Description	# Stays
0	11061003	Psychoactive substance use disorder (disorder)	2196
1	111479008	Organic mental disorder (disorder)	2082
2	11387009	Psychoactive substance-induced organic mental ...	982
3	15167005	Alcohol abuse (disorder)	1323
4	191480000	Alcohol withdrawal syndrome (disorder)	946
5	191492000	Drug-induced delirium (disorder)	540
6	191816009	Drug dependence (disorder)	1777
7	19943007	Cirrhosis of liver (disorder)	1988
8	235856003	Disorder of liver (disorder)	5203
9	2403008	Psychoactive substance dependence (disorder)	1323
10	243978007	Liver damage (disorder)	1115
11	268645007	Nondependent alcohol abuse (disorder)	1324
12	288293001	Poisoning by CNS drugs (disorder)	676
13	29212009	Alcohol-induced organic mental disorder (disor...	982
14	363101005	Drug withdrawal (disorder)	1088
15	41309000	Alcoholic liver damage (disorder)	1119
16	420054005	Alcoholic cirrhosis (disorder)	951
17	425413006	Drug-induced cirrhosis of liver (disorder)	951
18	427399008	Drug-induced disorder of liver (disorder)	1115
19	442351006	Mental disorder due to drug (disorder)	2502
20	7200002	Alcoholism (disorder)	1323
21	721710005	Fibrosis of liver caused by alcohol (disorder)	951
22	75478009	Poisoning (disorder)	872
23	7895008	Poisoning by drug AND/OR medicinal substance (...)	832
24	87858002	Drug-related disorder (disorder)	5017
25	91388009	Psychoactive substance abuse (disorder)	1767

Table 7: Medical concepts of heart disease graph. There are 3210 and 24150 unique stays of core nodes and entire graph respectively. Highlighted rows are the core nodes.

No.	SNOMED	Description	Stays
0	105981003	Disorder of cardiac function (disorder)	11610
1	11851006	Mitral valve disorder (disorder)	3212
2	123947004	Post-streptococcal disorder (disorder)	515
3	128599005	Structural disorder of heart (disorder)	13720
4	194732001	Diseases of mitral and aortic valves (disorder)	731
5	195002007	Multiple valve disease (disorder)	730
6	20721001	Tricuspid valve disorder (disorder)	773
7	23685000	Rheumatic heart disease (disorder)	521
8	301098007	Heart valve finding (finding)	5630
9	301099004	Aortic valve finding (finding)	2965
10	301101006	Mitral valve finding (finding)	3210
11	301108000	Tricuspid valve finding (finding)	770
12	368009	Heart valve disorder (disorder)	5630
13	40445007	Heart valve regurgitation (disorder)	702
14	415992005	Disorder of right cardiac ventricle (disorder)	1071
15	42343007	Congestive heart failure (disorder)	10983
16	448743001	Abnormality of aortic valve (disorder)	928
17	472762000	Disorder of right atrium (disorder)	1606
18	48724000	Mitral valve regurgitation (disorder)	681
19	56265001	Heart disease (disorder)	24245
20	8722008	Aortic valve disorder (disorder)	3099