

# Evaluating the Predictive Features of Person-Centric Knowledge Graph Embeddings: Unfolding Ablation Studies

Christos THEODOROPOULOS<sup>a,1</sup>, Natasha MULLIGAN<sup>b</sup> and

Joao BETTENCOURT-SILVA<sup>b,2</sup>

<sup>a</sup>*KU Leuven, Leuven, Belgium*

<sup>b</sup>*IBM Research Europe, Dublin, Ireland*

ORCID ID: Christos Theodoropoulos <https://orcid.org/0000-0002-2810-7817>

Joao Bettencourt-Silva <https://orcid.org/0000-0002-4562-0911>

Natalia Mulligan <https://orcid.org/0009-0006-5631-0325>

**Abstract.** Developing novel predictive models with complex biomedical information is challenging due to various idiosyncrasies related to heterogeneity, standardization or sparseness of the data. We previously introduced a person-centric ontology to organize information about individual patients, and a representation learning framework to extract person-centric knowledge graphs (PKGs) and to train Graph Neural Networks (GNNs). In this paper, we propose a systematic approach to examine the results of GNN models trained with both structured and unstructured information from the MIMIC-III dataset. Through ablation studies on different clinical, demographic, and social data, we show the robustness of this approach in identifying predictive features in PKGs for the task of readmission prediction.

**Keywords.** Entity-Centric Knowledge Graphs, Graph Neural Networks, Ablation Study, Patient Representation Learning, Hospital Readmissions.

## 1. Introduction and Background

Arguably the most sought-after predictive models in healthcare are those aimed at predicting readmissions. Avoiding subsequent, unplanned, hospitalizations has been a longstanding multidisciplinary effort involving clinicians, informaticians, researchers, statisticians as well as policy makers. Healthcare systems have introduced tracking metrics and incentives, with particular emphasis on conditions that are more prone to readmissions such as cardiovascular events or pneumonia [1]. Indeed nearly 1 in 4 heart failure (HF) patients are readmitted within 30 days, and approximately half of HF patients are readmitted within 6 months. Globally, readmissions are on the increase and reducing 30-day readmissions has now been a longstanding target to improve the quality of care and reduce expenditure [1-3]. Identifying the most relevant predictors has been reportedly challenging, especially for conditions such as HF [4, 5]. Recent efforts have sought to bring additional variables, parameters, and modalities to better understand the key drivers of readmissions. For example, it has been shown that notes written by physicians may have better capabilities for readmission prediction when compared with other modalities [4]. Other efforts have focused on creating a richer picture of patients

---

<sup>1</sup> Corresponding Author: Christos Theodoropoulos; E-mail: [christos.theodoropoulos@kuleuven.be](mailto:christos.theodoropoulos@kuleuven.be).

<sup>2</sup> Corresponding Author: Joao Bettencourt-Silva; E-mail: [jbettencourt@ie.ibm.com](mailto:jbettencourt@ie.ibm.com).

by including sociodemographic features, such as social determinants of health (SDoH), which are often embedded in clinical notes [6]. Finally, a recent approach focused on using neural networks and process mining for detecting important predictors based on time information associated with hospital events, severity scores, and demographics [7].

Graph Neural Networks (GNNs) [5, 8] have emerged as powerful tools for learning representations from graph-structured data from various domains. By leveraging the inherent structure and connectivity of graphs, GNNs offer a versatile approach for capturing complex relationships. However, methodological frameworks to continuously evaluate model results with respect to a given prediction tasks are needed, especially in healthcare domains. This paper proposes a method to systematically evaluate expert-driven hypotheses about predictors of interest and their impact on GNN model results. Our approach centers on learning representations for person-centric graphs using GNNs, and, by encoding the intricate interplay between patients' clinical attributes, demographics, and social aspects within the graph structure, we have developed GNN-based predictive models capable of identifying patients at high risk of readmission.

## 2. Methods

We previously developed an ontology that defines person-centric knowledge graphs (PKG) and their social context, and studied how different ontology-driven graph representations and heterogeneity impact models' performance [5]. In this paper, we introduce a repeatable method and approach to continuously identify predictive features of interest based on ablation studies (Figure 1A). The approach requires as input: a task (e.g., readmission prediction), a dataset (e.g., hospital health records) and a person-centric graph schema (e.g., ontology defining the available data points and how they relate to a central node of interest). Before training a GNN, the following pre-processing steps are required to create PKGs: (1) data must first be selected from the input dataset, (2) an assessment of the data and its quality (e.g. missing data) is performed, (3) data can further be enriched by mining unstructured text notes or introducing other modalities, (4) the data is then sampled and (5) a summary of the dataset is created so that knowledge graphs can be extracted. Our framework for GNN model training requires individual graphs to be created for each patient hospital admission. Once the model is trained, an extensive ablation study is conducted to unearth the most predictive features encoded within the graph representation. Finally, an analysis and interpretation of model results is carried out to dissect the contributions of different features, such as demographic variables and social determinants of health, in the predictive performance of the model.

### 2.1. Pre-processing and Graph Neural Network Model Training

The selected task was predicting Intensive Care Unit (ICU) readmissions using the MIMIC-III dataset<sup>3</sup>. In this paper, we rely on a representation learning framework previously developed [6] and the Health & Social Patient-centric Ontology (HSPO) to define how PKGs are created. Readmissions are defined as any subsequent admissions within 30 days of the first ICU visit in MIMIC-III. Data was extracted such that PKGs contain information about demographic, clinical, and social context, each represented by distinct graph nodes. The demographic nodes encapsulate attributes such as age group, gender, religion, marital status, and race. For the clinical view, nodes represent diseases,

---

<sup>3</sup> MIMIC-III Clinical Database. 2016. <https://physionet.org/content/mimiciii/1.4>

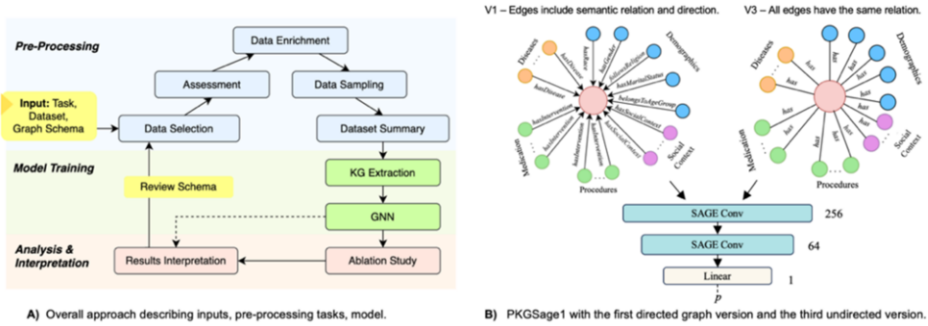


Figure 1. Overall approach (A), Knowledge Graphs and GNN (B).

received medication, and undergoing procedures. Additionally, to incorporate social aspects, information regarding employment status, household composition, and housing conditions is extracted from clinical notes using UMLS codes sampled via the MetaMap annotator<sup>4</sup> and represented as nodes in the graphs.

We utilize the best performing model PKGSage<sub>1</sub> [5] that leverages the Sage Graph Convolution Network architecture (Figure 1B). Determining the most effective graph structure is a challenging task and was previously addressed in [5]. We conducted experiments using the directed first graph version and the undirected third version (Figure 1B), both of which demonstrate the best performance.

## 2.2. Ablation Study, Analysis and Interpretation

The ablation study aims at comprehensively assessing the predictive ability of different features and facets (i.e. group of features) within the person-centric graph embeddings. The objective is to systematically evaluate the impact of excluding different information on the model's predictive performance. Initially, each facet is removed, and the resulting performance degradation is assessed. To obtain further insights, different combinations of features (e.g. all medications, procedures) are also excluded. Given the complex interdependencies between clinical information, we explore the robustness of the models by conducting experiments at a second level, where two facets of the clinical view are entirely excluded, followed by the exclusion of all clinical facets (diseases, medication, and procedures). Furthermore, considering the ICU readmission prediction task, further experiments are conducted with a focus on excluding the most well-known and frequent conditions in U.S. hospitals for all-cause 30-day readmissions<sup>5</sup>.

## 3. Results and Discussion

The results of the ablation study (Figure 2A, B) shed light on the importance of demographic information, particularly in the context of readmission prediction models. We observed that the removal of socio-demographic variables impacts model performance, particularly in terms of accuracy. Among these, marital status and race emerged as the two strongest predictors within our dataset. It is noteworthy that marital status, in addition to its demographic significance, also carries a social dimension

<sup>4</sup> MetaMap. 2016. [https://www.nlm.nih.gov/research/umls/implementation\\_resources/metamap.html](https://www.nlm.nih.gov/research/umls/implementation_resources/metamap.html)

<sup>5</sup> Statistical Brief #153. HCUP. May 2016. US AHRQ. <https://hcup-us.ahrq.gov/reports/statbriefs/sb153.jsp>

Information Excluded	PKGSage, undirected V3		PKGSage, directed V1		Information Excluded	PKGSage, undirected V3		PKGSage, directed V1	
	Accuracy	F1	Accuracy	F1		Accuracy	F1-Score	Accuracy	F1-Score
-	61.69	68.06	62.16	67.93	-	61.69	68.06	62.16	67.93
Demographics	60.42 (-2.06%)	67.4 (-0.97%)	60.46 (-2.73%)	67.34 (-0.87%)	Diagnoses (D)	<b>59.69 (-3.24%)</b>	<b>66.87 (-1.75%)</b>	<b>59.59 (-4.13%)</b>	67.48 (-0.66%)
Age	60.52 (-1.9%)	67.94 (-0.18%)	60.69 (-2.36%)	<b>67.21 (-1.06%)</b>	Readmission				
Gender	60.46 (-1.99%)	67.53 (-0.78%)	60.49 (-2.69%)	67.32 (-0.9%)	Diagnoses (RD)	60.23 (-2.37%)	67.34 (-1.06%)	60.53 (-2.62%)	67.33 (-0.88%)
Religion	60.49 (-1.95%)	68.01 (-0.07%)	60.51 (-2.65%)	67.58 (-0.52%)	Medications (M)	59.7 (-3.23%)	66.98 (-1.59%)	60.03 (-3.43%)	<b>66.8 (-1.66%)</b>
Marital Status	60.31 (-2.24%)	<b>67.26 (-1.18%)</b>	<b>60.23 (-3.1%)</b>	67.56 (-0.54%)	Procedures (P)	60.42 (-2.06%)	67.43 (-0.93%)	59.97 (-3.52%)	67.56 (-0.54%)
Race	<b>60.2 (-2.42%)</b>	67.74 (-0.47%)	60.3 (-2.99%)	67.37 (-0.82%)	D, M	59.49 (-3.57%)	<b>65.69 (-3.48%)</b>	59.72 (-3.93%)	<b>66.06 (-2.75%)</b>
Social Aspects	60.68 (-1.64%)	67.63 (-0.63%)	60.5 (-2.67%)	67.64 (-0.43%)	D, P	59.32 (-3.84%)	66.62 (-2.12%)	59.53 (-4.23%)	66.38 (-2.28%)
Employment	60.65 (-1.69%)	67.84 (-0.32%)	60.46 (-2.73%)	<b>67.36 (-0.84%)</b>	M, P	<b>58.76 (-4.75%)</b>	66.59 (-2.16%)	<b>59.47 (-4.33%)</b>	66.86 (-1.58%)
Housing	60.23 (-2.37%)	67.5 (-0.82%)	60.88 (-2.06%)	67.5 (-0.63%)	D, P, M	<b>51.84 (-15.97%)</b>	<b>62.36 (-8.37%)</b>	<b>50.82 (-18.24%)</b>	<b>64.32 (-5.31%)</b>
Household	<b>60.17 (-2.46%)</b>	<b>67.3 (-1.12%)</b>	<b>60.3 (-2.99%)</b>	67.77 (-0.24%)	RD, P, M	59.4 (-3.71%)	66.55 (-2.22%)	58.28 (-6.24%)	65.37 (-3.77%)

A) Ablation study results: Demographic and Social information.

B) Ablation study results: Diagnoses (D), known Readmission Diagnoses (RD), Medications (M), Procedures (P).

Figure 2. Ablation study results (A, B).

pertaining to the patient's support network and living arrangements. Contrary to initial expectations, age did not emerge as the strongest predictor in our analysis. However, a deeper examination of the dataset revealed that a significant proportion of patients are elderly, with 56.24% being over 70. This underscores the unique demographic composition of the dataset and emphasizes the need for a nuanced analysis.

Despite the large number of missing social information, our study demonstrated that excluding social aspects detrimentally affects model performance. The exclusion of household information (77.9% missing data) had the most performance impact. This finding underscores the critical role of social determinants in predictive models and the importance of robust data collection strategies. While the systematic acquisition of social data poses challenges, such as dealing with unstructured data [9], extracting knowledge from biomedical text [10], and defining a standardized protocol, it is imperative for enhancing a data-driven understanding of readmission prediction or other tasks.

Regarding clinical predictors (Figure 2B), we find that the pure clinical view of patients emerges as a key driver, with diseases and medications standing out as the most influential factors. One notable observation is the inherent inter-correlations among various aspects of the clinical view (e.g., patients diagnosed with specific diseases will receive disease-specific medications and procedures). While excluding the most common diagnoses associated with high readmission rates has a discernible effect on model performance, the exclusion of the entire disease information leads to an additional drop in performance metrics, up to 4.13% percentage decrease in accuracy. This suggests that additional diagnoses beyond the most known ones may play a substantial role in driving ICU readmissions. These findings underscore the importance of expanding the scope of diagnoses (and ICD code mappings) considered in predictive modeling efforts to capture a more comprehensive view of readmission risk factors. Medication and procedures are also strong predictors (Figure 2B).

Regarding the exclusion of paired clinical information, such as disease-medication or medication-procedure combinations, we observe a further impact on model efficiency, albeit not as significant as the exclusion of entire facets. This reaffirms the inter-correlation among clinical factors and underscores the need for holistic modeling approaches. Excluding clinical information completely results in significant performance degradation with up to 18.24% and 8.37% percentage decline in accuracy and F1 score (i.e., the harmonic mean of precision P and recall R:  $F1 = 2 \frac{(P \cdot R)}{(P + R)}$ ) respectively. Hence, this supports the fact that the clinical view of the patient may be the best predictor for the ICU readmission prediction task given the available data.

Overall, we highlight that the resilience of the models in managing missing data is profound, as the decline in performance is marginal when we exclude one facet of the

data. In the worst case, accuracy and F1 score present 4.13% and 1.75% percentage decrease respectively. The models exhibit robust performance even when two of the three clinical facets are not present, with up 4.75% and 3.48% percentage deterioration in accuracy and F1 score correspondingly. We emphasize that stability and resilience are crucial characteristics, particularly in healthcare datasets where missing data is unavoidable due to privacy concerns, system limitations, or human errors.

#### 4. Conclusions

This paper presents a novel and data-driven approach to identify GNN model predictors through conducting ablation studies within the framework of patient-centric graph prediction tasks. Through an extensive analysis focused on the ICU readmission prediction task, we identify and highlight the most influential predictors, shedding light on critical factors that impact patient outcomes in clinical settings. One of the key takeaways is the paramount importance of systematic collection of social determinants data, as we detect performance degradation even if this information availability is largely missing in the MIMIC-III dataset. Our framework provides flexible and adaptable methodological steps that can be applied to different datasets and predictive tasks within the healthcare domain and beyond. These steps allow the continuous examination of new predictors, which can be domain expert-driven, or as future work, using generative models. Future work will also extend the application of our framework to diverse healthcare scenarios and hypotheses. By leveraging the insights from the ablation studies, we aim to support the advancement of robust predictive models that can be interpretable and actionable, ultimately contributing to improved care and service delivery.

#### References

- [1] Khan MS, et al. Trends in 30- and 90-Day Readmission Rates for Heart Failure. *Circ Heart Fail.* 2021 Apr;14(4):e008335. doi: 10.1161/CIRCHEARTFAILURE.121.008335. Epub 2021 Apr 19. PMID: 33866827.
- [2] Lawson C, et al. Trends in 30-day readmissions following hospitalisation for heart failure by sex, socioeconomic status and ethnicity. *EclinicalMedicine.* 2021 Jul 14;38:101008. doi: 10.1016/j.eclinm.2021.101008. PMID: 34308315; PMCID: PMC8283308.
- [3] Meneill H, Khairat S. Impact of Intensive Care Unit Readmissions on Patient Outcomes and the Evaluation of the National Early Warning Score to Prevent Readmissions: Literature Review. *JMIR Perioper Med.* 2020 May 8;3(1):e13782. doi: 10.2196/13782. PMID: 33393911; PMCID: PMC7709858.
- [4] Sheerit E, et al. Predicting unplanned readmissions in the intensive care unit: a multimodality evaluation. *Sci Rep.* 2023 Sep 18;13(1):15426. doi: 10.1038/s41598-023-42372-y.
- [5] Theodoropoulos C, et al. Representation Learning for Person or Entity-Centric Knowledge Graphs: An Application in Healthcare. *Knowledge Capture Conference.* 2023 Dec 5. doi: 10.1145/3587259.3627545.
- [6] Bettencourt-Silva JH, et al. Discovering New Social Determinants of Health Concepts from Unstructured Data: Framework and Evaluation. *Stud Health Technol Inform.* 2020 Jun 16;270:173-177. doi: 10.3233/SHTI200145. PMID: 32570369.
- [7] Pishgar M, et al. Prediction of unplanned 30-day readmission for ICU patients with heart failure. *BMC Med Inform Decis Mak.* 2022 May 2;22(1):117. doi: 10.1186/s12911-022-01857-y. PMID: 35501789; PMCID: PMC9063206.
- [8] Wu Z, et al. A Comprehensive Survey on Graph Neural Networks. *IEEE Trans Neural Netw Learn Syst.* 2021 Jan;32(1):4-24. doi: 10.1109/TNNLS.2020.2978386. Epub 2021 Jan 4. PMID: 32217482.
- [9] Theodoropoulos C, et al. Imposing Relation Structure in Language-Model Embeddings Using Contrastive Learning. *Conference on Computational Natural Language Learning.* 2021 Nov. doi: 10.18653/v1/2021.conll-1.27.
- [10] Theodoropoulos C, Moens MF. An Information Extraction Study: Take in Mind the Tokenization!. *Conference of the European Society for Fuzzy Logic and Technology.* 2023 Aug. doi: 10.1007/978-3-031-39965-7\_49.