Contents lists available at ScienceDirect
Big Data Research
www.elsevier.com/locate/bdr

Predicting Adverse Events After Surgery ☆

Senjuti Basu Roy [a], Moushumi Maria [b], Tina Wang [b], Anne Ehlers [c], David Flum [c]

[a] CS, NJIT, United States of America
[b] UW Tacoma, United States of America
[c] Dept of Surgery, UW Seattle, United States of America

ARTICLE INFO

Article history:
Received 24 August 2017
Received in revised form 15 February 2018
Accepted 19 March 2018
Available online xxxx

ABSTRACT

Predicting risk of adverse events (AEs) following surgical procedure is of significant interest, as that may guide in better resource utilization and an improved quality of care. Currently available comorbidity indices are largely inaccurate to predict adverse events other than death, as well as off-the-shelf machine learning models do not typically account for the temporal sequence of events to enable predictive analytics. We propose a study to improve the current techniques for assessing and predicting the risk of adverse events (AEs) associated with multiple chronic conditions by designing machine learning models that account for and incorporate the temporal sequence and timing of conditions. We formalize the task as a binary classification problem. Our technical contributions include *devising novel sequence based feature discovery techniques to augment existing supervised classification algorithms, as well as formalizing the classification task as a Markov Chain Model (MCM) that captures the temporal sequence of prior chronic conditions/events*. Finally, we design a *hybrid or multi-classifier that combines prediction from the aforementioned classification models to finally predict AE*. Our experimental results, conducted using the Truven Health MarketScan Research Databases with more than 27 million of claim records on two different surgery types, discover interesting insights that can guide patient-centered decision-making and can direct healthcare teams to adjust techniques and interventions. We also extensively compare the performance of our solutions to appropriate baselines.

1. Introduction

More than 80 million surgeries are performed annually in the United States, increasingly among patients with one or more chronic conditions. Surgical procedures [1,2] are increasingly being performed in those with multiple chronic conditions, and these patients have the greatest risk of serious complications and procedure-related deaths. Although complication rates from major surgery rise with age and multiple conditions, the effect of specific combinations or sequences of comorbid conditions [3,4] on outcome is not well understood. Identifying patients at increased risk for complications and adverse outcomes is critical to direct healthcare teams to adjust techniques or interventions, improve decision-making and quality improvement.

The most common technique used to quantify the burden of health conditions using secondary data are comorbidity indices, or numeric scores based on a summation of the number of conditions [3,4] that apply pre-determined "weights" to give certain conditions more importance over others. Comorbidity indices [3,4], however, are subject to several limitations and have not been widely incorporated into the routine assessment of patients in clinical care.

The last decade has also seen a rapid increase in the number of available techniques for building predictive models [5–13] for different applications with much larger numbers of attributes, such as predicting risk of readmission, risk of emergency visit post discharge, risk of dementia, etc. However, most of these past efforts do not account for the temporal sequence of events/chronic conditions to enable predictive analytics considering only large scale claim data. Rather, they focus on capturing the prior events independently to build supervised classification models mostly considering clinical attributes.

We hypothesize that the current techniques for assessing and predicting the relationship of multiple health conditions and outcome may be improved in several ways: 1) include condition specific diagnoses and condition specific outcomes; 2) evaluate specific combinations of conditions to assess if their contribution to the risk of outcome is something other than additive; 3) determine if the temporal sequence of chronic conditions contributes to the prediction of risk beyond the conventional assessment of

☆ This article belongs to Special Issue: Medical Data Analytics.

E-mail address: senjutib@njit.edu (S.B. Roy).

https://doi.org/10.1016/j.bdr.2018.03.003
2214-5796/

Table 1
Claim history of the patient in Example 1.

Claim-no	Date	Conditions	AEs	Surgery	Age	Sex	Region
1	12-12-2010	pvd (ER)	no	no	54	male	XYZ
2	12-25-2010	mi (In)	no	no	54	male	XYZ
3	2-9-2011	chf (ER)	cardiac disease	no	54	male	XYZ
4	5-5-2011	no	no	large bowel (In)	54	male	XYZ
5	5-29-2011	mi (Out)	cardiac disease	no	54	male	XYZ
6	6-15-2011	no	sepsis (ER)	no	54	male	XYZ

whether the conditions are present at all. We formalize the task as a binary classification problem and present novel machine learning models that *account for and incorporate the temporal sequence and timing of conditions*. Our technical contributions include *devising novel sequence based feature discovery techniques to augment existing supervised classification algorithms, as well as formalizing the classification task as a Markov Chain Model (MCM) [14] that captures the temporal sequence of chronic conditions. We then design a* hybrid classifier or multi-classifier *[15] that combines the MCM classifier with other classification models to generate the final prediction.*

We compare our proposed solutions with existing index based risk stratification techniques [3,4], as well as with classification algorithm that does not leverage temporal sequence. We use the Truven Health MarketScan Research Databases [16] with more than 27 million of records on two different surgery types for empirical analysis. Our analysis discovers interesting insights that can guide patient-centered decision-making, and can direct healthcare teams to adjust techniques and interventions, as well as exhibits significant performance improvement compared to the baseline methods.

In summary, the paper makes the following contributions:

- We formalize the task of predicting AE after surgery with one or more chronic conditions as a binary classification problem (Section 1.1).
- We first propose novel feature discovery methods that *account for and incorporate the temporal sequence and timing of chronic conditions* and build effective classification models using them. Additionally, we formalize the classification problem using a Markov Chain Model (MCM) that captures the temporal sequence of chronic conditions. We finally design a hybrid classifier that incorporates the judgments from the individual classifiers to predict AE (Section 2).
- We conduct large scale experiments using the *Truven Health MarketScan Research Databases [16] with more than 27 million of claim records on two different surgery types that discover interesting insights that can guide patient-centered decision-making and can direct healthcare teams to adjust techniques and interventions (Section 3.5).*
- Our experimental results exhibit *statistically significant performance improvement compared to multiple baseline methods* (Section 3.6).

1.1. Problem formulation

The Truven Health MarketScan Databases contain claim information, where each claim is associated with a patient identifier (id) and a date, as well as socio-demographic information such as age, sex, plan type, and region of the United States. More importantly, each claim has information about the patient comorbidities or chronic conditions, such as diabetes, congestive heart failure (CHF), or cancer. The details of the chronic conditions (identified using International Statistical Classification of Diseases 9 (ICD-9) codes) could be found in Table 3. Our domain expert collaborators define – reoperation, sepsis, cardiac arrest, and death as some of the AEs. The details of the AEs used in our datasets can be found in Table 4. Additionally, the data contains information of the encounter type of the patients – whether she was treated in the hospital (Inpatient/In), seen by a physician at an outpatient clinic (Outpatient/Out), or was serviced in an emergency room (ER).

With every claim, we get to know whether a surgery occurred or not using Current Procedural Terminology (CPT) codes. A patient may have multiple claims associated with different dates, some of these claims could be for the surgical intervention, while the rests are only for treating chronic conditions or adverse events before or after surgery. If a patient has undergone multiple surgeries, we will consider each surgery individually to predict adverse event associated with it.

By considering patients' healthcare utilization and claims in the six months prior to undergoing surgery, our task is to predict the likelihood of encountering an AE within three months of surgery. While one cannot ever be 100% confident that any particular event or death is associated with surgery, this timeline is commonly used by the policy makers and payers[1] to decide whether or not an AE is associated with the surgery itself.

Example 1. We now present the claim history of a 54 old male who has undergone a Large Bowel surgery in the state of XYZ. The patient has 3 claims associated with different types of chronic conditions before surgery, some of which also had adverse events associated. He has an additional claim on date (5-5-2011) that is only associated with surgery. Within 90 days after the surgery, the patient has two additional claims with adverse events. For every claim, his age and region information is also present. Table 1 contains the claim history of this patient.

We simplify the problem further as a *binary classification task* with the objective to predict if the patient will have any of the adverse events (AEs) within 90 days after surgery. This way, for every patient in training or testing, the class label is the AE indicator flag which is set to 0 or 1 based on this simplification. Using Example 1, the patient had two adverse events within 90 days after surgery, hence the class label associated with this patient is 1.

Problem 1 *(AE prediction problem).* The feature vector $X_i = (x_{i1}, x_{i2}, \ldots, x_{iM})$ of a patient i includes information about general demographics such as age and gender, as well as specific chronic conditions that she has encountered *six months prior to surgery*. The goal is to predict a binary class label AE_i ($AE_i = 0$ means no adverse event *within 90 days of surgery* and $AE_i = 1$ means otherwise) for every test instance. Given training examples of the form (X_i, AE_i) with $X_i \in X$ and $AE_i \in \mathcal{AE}$, the objective is to learn a model $\mathcal{H} : X \to \mathcal{AE}$ that can label new, unseen instances from X with a label from $\mathcal{AE}$ in an accurate way.

2. Predicting adverse events

Consider Example 1 and note that in order to successfully predict AE after surgery, one has to look for the *temporal sequence of chronic conditions/events that have taken place six months before*

[1] https://www.medicare.gov/hospitalcompare/Data/30-day-measures.html.