Siegfried Horschig, Trends in Bioinformatics WS 2017/2018

Prediction of Patient Outcomes after Renal Replacement Therapy in the ICU

Motivation

The renal system in the human body has the purpose to eliminate wastes from the body and control levels of certain substances in the blood. If this system is impaired, for example due to Acute Kidney Injury (AKI), artificial methods in the form of Renal Replacement Therapy (RRT) have to be introduced, more commonly known under the term of dialysis. Dialysis' outcomes are highly dependent on the patient's characteristics. To support clinical decision making, our goal is to develop a Clinical Prediction Model (CPM), which predicts patient outcomes based on their data while in the Intensive Care Unit (ICU). Additionally, we want the resulting model to be interpretable. Especially in the medical context, it is important to know why a specific decision was made to ensure patient safety and validate that decision. More powerful models, such as Deep Neural Networks, do not expose their decision process in a human-readable fashion and are thus non-interpretable. Our goal was also to make such a non-interpretable model interpretable.

Related Work

We are building on top of existing studies researching parameters of RRT such as [1] or [2]. Different models for this prediction tasks have also been discussed in [3] or [4]. The interpretability of machine learning models is further explored in [5] and [6], while the latter also introduces mimic learning as a technique.

Methods

The MIMIC III database provides the data for our model. This includes, but is not limited to, blood values such as creatinine, patient demographics or the glomerular filtration rate. As target outputs of the model we define various parameters to estimate how healthy the patient is, such as the length of his ICU stay, the number of days he spends ventilation free or his mortality rate.

As interpretable model to train on this data, we used an algorithm called Bayesian Rule Lists (BRL) [7]. As non-interpretable counterpart, our choice was a Deep Neural Network, realized by the scikit-learn library with the Multi-Layer-Perceptron Classifier (MLP) [8]. The method we incorporated to make the predictions of the MLP explainable is called Mimic Learning. We do this by first training the MLP with the required input data. Afterwards, we train our BRL classifier with the predicted probabilities of the MLP to best mimic its behaviour. In the best case, we end up with a model that behaves like the MLP, but with an explainable structure lying beneath.

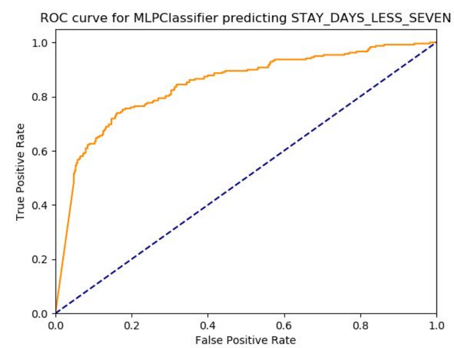

Fig. 1: ROC performance of the MLP classifier when predicting if the length of the ICU stay exceeds 7 days

Preliminary Results

	Recall	Specificity	Precision
MLP Classifier	63%	76%	65%
Bayesian Rule Lists	66%	67%	58%

Table 1: Comparison of the two different models in the prediction of mortality.

As seen in Table 1 and Fig. 1, results show us that the MLP outperforms our BRL classifier as expected, but by a rather small margin. Difficulties arise when trying to incorporate this model in the mimic learning process, as the existing implementation only supports binary classification. Thus, we can't use the predicted probabilities to train our mimic model and when training on the absolute predictions, we also let the BRL model learn the errors of the MLP.

But even though the performance is in need of improvement, the results show some insight into the more important features determining the outcome of this classification. In Fig. 2, we can see an excerpt of the output the BRL classifier gives us.

IF GFR_48_B : -inf_to_5.9937927112524605 AND ELIXHAUSER_VANWALRAVEN : -inf_to_11.5 THEN
probability of DIED_90DAYS: **3.7%** (1.9%-6.2%)
ELSE IF LACTATE : 4.235_to_inf AND BICARBONATE : -inf_to_21.925 THEN
probability of DIED_90DAYS: **77.4%** (71.2%-83.0%)
ELSE IF LENGTH_OF_STAY_HOURS : 1030.5_to_1475.5 AND OBESITY : All THEN
probability of DIED_90DAYS: **98.5%** (95.9%-99.8%)

Fig. 2: Example output of the Bayesian Rule List classifier. Different values and their influence can be seen, such as the Glomerular Filtration Rate 48 hours before the procedure began or the time in hours the patient spent in the ICU.

[1] "Dose of Dialysis in Acute Renal Failure." http://cjasn.asnjournals.org/content/1/3/380.full.
[2] "Renal replacement therapy for acute kidney injury" 27 Jun. 2016, https://rrtjournal.biomedcentral.com/articles/10.1186/s41100-016-0043-1.
[3] "Performance Comparison of Three Data Mining Techniques for Predicting Kidney Dialysis" http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=26E33F771F4135FBD2451A2768D0492D?doi=10.1.1.668.9133&rep=rep1&type=pdf.
[4] "Predicting survival time for kidney dialysis patients: A data mining approach" https://www.ncbi.nlm.nih.gov/pubmed/15749092.
[5] Machine Learning Model Interpretability for Precision Medicine." 28 Oct. 2016, https://arxiv.org/abs/1610.09045
[6] "Interpretable Deep Models for ICU Outcome Prediction" https://www.ncbi.nlm.nih.gov/pubmed/28269832.
[7] "Scalable Bayesian Rule Lists" 27 Feb. 2016, https://arxiv.org/abs/1602.08610.
[8] "sklearn.neural_network.MLPClassifier" http://scikit-learn.org/stable/modules/generated/sklearn.neural_network.MLPClassifier.html.