Early detection of acute kidney injury with Bayesian networks

Harry Cruz[1][2], Bastien Grasnick[1], Henriette Dinger [1], Frank Bier[2], and Christoph Meinel[1]

[1]Hasso Plattner Institute, Potsdam

[2]Fraunhofer Institute for Cell Therapy and Immunology, Potsdam

`Harry.FreitasDaCruz@hpi.de` (corresponding author)

Abstract

Acute kidney injury (AKI) is a major health issue, affecting large numbers of patients worldwide. It is associated with an increase in complications and poor prognostics if diagnosis is delayed. Medical guidelines are routinely employed to classify different AKI stages, but guidance on the early detection of AKI risk is limited. In this paper, we present a Bayesian Networks (BN) proof of concept to predict the likelihood of AKI onset based on longitudinal patient data, such as serum creatinine values, demographics and comorbidities. Data for training and validating the model was obtained from the Multiparameter Intelligent Monitoring in Intensive Care (MIMIC II) database. We describe the problem domain, data acquisition and preparation, model developed, results obtained and pertaining limitations. We demonstrate that our model can predict the onset of the disease with an accuracy of up to 87% (area under the curve of 0.87) in the cohort under analysis.

1 Introduction

Acute Kidney Injury (AKI) affects a large portion of the elderly population and has a high risk of death, as there is no trivial treatment once it breaks out (Statistisches Bundesamt, 2014). After onset, the patient may even need dialysis and/or renal replacement therapy (kidney transplant). Currently, detection of kidney injury requires continuous monitoring of creatinine and other lab values (Harty, 2014). In particular, when many patients must be monitored at once, it is hard for physicians to keep track of subtle changes in blood measurements which might be indicative of AKI. As a consequence, a significant portion of patients is diagnosed for AKI too late, leading to more complications and higher mortality. In fact, a study in the UK found that 60% of post-admission AKI cases were avoidable (Stewart et al., 2009). An automated, early detection of high-risk patients may lead to a faster response of physicians, reducing complications associated with AKI.

Our objective was therefore to develop a proof of concept for early detection of AKI. For this purpose, we created and trained a Bayesian network model on the basis of real patient data. A Bayesian network is a probabilistic graphical model, consisting of random variables and their influences on one another. Data for training and validating the model was obtained from the anonymized Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC II) database (Lehman et al., 2011). In this paper, we will provide the background needed and present the methods involved in developing the model, including data acquisition and preparation, results obtained and further discussion.

2 Background

This section deals with the necessary background for the remainder of this paper. This includes an elucidation of risk factors related to AKI, Bayesian networks fundamentals as well as related work.

2.1 Risk factors for AKI

As a starting point, we needed to identify factors which predispose patients to AKI from medical literature sources. These are, among others, creatinine values taken from blood or urine samples. Furthermore, comorbidities, such as heart failure or diabetes and personal background, including age, gender and ethnicity, have to be considered. Since we are in the domain of kidney diseases, dehydration plays an important role too (Lopes and Jorge, 2013; Kellum et al., 2012). In detail, the relevant factors are:

- **laboratory values** (serum creatinine, urine output, estimated gloremular filtration rate (eGFR) value)

- **comorbidities** (heart failure, chronic kidney disease, tumor disease, diabetes, obesity, hypothyroidism, paralysis, hypertension, pulmonary circulation, valvular disease, peptic ulcer, deficiency anemia, renal failure)

- **personal background** (age, gender, ethnicity, admission type, that is emergency or elective)

2.2 Acute Kidney Injury classification

Currently, two main guidelines are used in medicine for the classification of AKI: RIFLE (Risk, Injury, Failure, Loss) and AKIN (Acute Kidney Injury Network). Both help physicians establish severity of kidney injury based on the serum creatinine and urine output of a patient. Figures 2 and 1 (Cruz et al., 2009) show an overview of the two classifications depending on the creatinine and urine values.

Figure 1: **RIFLE classification for AKI** The RIFLE classification uses serum creatinine and urine output values. It consists of five classes: risk, injury, failure, loss of kidney function and end-stage kidney disease.

The RIFLE classification (Bellomo et al., 2004; Ricci et al., 2011) predates AKIN. It consists of five stages: risk, injury, failure, loss of kidney function and end-stage kidney disease (ESKD). In comparison to that, the AKIN classification (Ricci et al., 2011; Mehta et al., 2007) uses only three stages: risk, injury, and failure (Lopes and Jorge, 2013; Kellum et al., 2012). Since AKIN is based on RIFLE, it is more widely used nowadays. AKIN performs better for detecting early stage

Figure 2: **AKIN classification for AKI** The classification system works with the help of serum creatinine and urine output values. There are three possible categories: risk, injury and failure.

patients, while RIFLE guideline is better suited for patients in advanced stage of renal function loss, while Since both guidelines are well-proven in practice, they will serve as an additional output variables for the model (Lopes and Jorge, 2013; Kellum et al., 2012).

2.3 Bayesian networks

Howard et al. define a Bayesian network as "an annotated directed graph that encodes probabilistic relationships among distinctions of interest in an uncertain-reasoning problem" (Howard and Matheson, 1983). In general, a Bayesian network consists of multiple random variables and their conditional dependencies modeled as probability functions. This way, based on evidence provided for one or more given variables, the probability of the other random variables can be calculated after the network was trained (Horný, 2014).

An example is given in Figure 3 adapted from (Barbini et al., 2013), where the probabilistic dependencies of a simple Bayesian model with four dichotomous variables (true or false) is shown. It follows from the model that A and B (having *a priori* associated probalities) exert influence on C. In turn, this effect is modeled by a conditional probability table on the children node. As such, the probability of event C occurring given that A occurred but not B is given by $X_{A\overline{B}}$. Nodes C and D are independent of each other (conditional independence).

A Bayesian network is developed either by using expert knowledge and building the network manually or letting the network be built directly

from the data by a specific algorithm, an approach known as structure learning. Once the structure of the network is learned, it can be further manipulated using expert knowledge. Moreover, when the structure is set, the probability functions representing the conditional dependencies can also be learned from data. This is referred to as parameter learning.

The most significant drawback of Bayesian networks is that the accuracy depends highly on the chosen structure of the model. If this is done negligently, the resulting model can fail to show existing results (false negatives) or show incorrect results (false positives). A common way to avoid this is to iteratively establish dependencies among variables, usually based on expert knowledge (Heckerman et al., 1995).

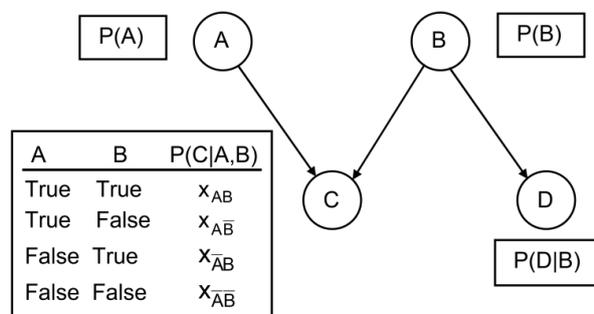| A	B	$P(C|A,B)$
True	True	x_{AB}
True	False	$x_{A\overline{B}}$
False	True	$x_{\overline{A}B}$
False	False	$x_{\overline{A}\overline{B}}$

Figure 3: **Example of a Bayesian network** This Bayesian network shows four random dichotomous variables and their probabilistic relationships.

2.4 Previous work

Machine learning has been widely utilized in the medical domain in several instances. Particularly in Nephrology, Legrand et al. evaluated the post-operative AKI risk of patients suffering from infective endocarditis after undergoing cardiac surgery. They applied super learning, a technique to choose the optimal regression algorithm, comparing ten different models by using cross-validation. Targeted maximum likelihood estimation was used to obtain the following most important risk factors: multiple surgery, pre-operative anemia as defined by a baseline hemoglobin level <10 g/dl, transfusion requirement during surgery, the use of a nephrotoxic agent: vancomycin, aminoglycoside or contrast iodine; and the interaction between vancomycin and aminoglycoside. (Legrand et al., 2013).

Further, Król et al. developed an approach to predict chronic kidney disease (CKD). They did not build a technical system but "an algorithm for the diagnostic procedure" (Król et al., 2009). For this purpose, they did an investigative survey with 2471 randomly chosen people involved. As a result, they found different factors that encourage a CKD. Among others, these factors are the male gender, diabetes and hypertension.

Specifically applying Bayesian networks, Onisko et al. present a model based on dynamic BNs for predicting the risk of cervical cancer, using hospital data and expert knowledge. The authors were able to categorize patients in different risk categories (Onisko et al., 2004). In a similar approach, Nachimuthu et al. used BNs for early detection of sepsis(Nachimuthu and Haug, 2012).

Similarly, Ward et al. offer a framework for the development of Bayesian networks in the particular example of sepsis. They build their model based on knowledge gained from literature, hospital data as well as expert knowledge. Their resulting model provides a base for a correct prediction. Since the data set is rather small, a further evaluation is planned to support their result (Ward et al., 2014).

In an approach analogous to these works, we developed a model based on Bayesian networks for estimating the risk of developing AKI. We were also able to use hospital data and an expert consultation for the development. To the best of the authors' knowledge, this is the first work explicitly utilizing a Bayesian network model for AKI prediction.

3 Model development

3.1 Methodology

For the model development, we utilized two machine learning tools, Weka and GeNIe and compared their accuracy to control for possible tool bias in the results. Further, we extracted the needed data from the MIMIC II database, which was preprocessed for tool input. We created two data sets for cross-validation, one with 6000 entries and another with 9000 entries (50% more). In the first iteration, the AKI literature laid out in section 2.1 formed the basis for the development of an initial model (1st iteration model). This model was then augmented and corrected after an expert consultation session with nephrologists at the Charité hospital in Berlin (2nd iteration model). We then

compared the two models, as well as the different tools and analyzed the results obtained. The following sections will provide further details into this procedure.

3.2 Tools utilized

In an effort to avoid bias in the results possibly introduced by differing algorithm implementations, we chose to develop and test the model in two widely available Bayesian network modelling tools, Weka and GeNIE.

Apache Weka is a Java toolkit for different kinds of data mining algorithms. It allows the classification, clustering and visualization of data sets (Kumar and Sahoo, 2012). One of the main advantages of Weka is the very powerful capabilities for Bayesian networks (Bouckaert, 2008). For network structure learning, an estimator as well as a search algorithm can be set as parameters in the tool. For the purposes of this paper, we chose the algorithm K2 as it has the best performance among the search algorithms implemented in Weka (The University of Waikato, 2008).

GeNIe is the user interface of SMILE, a C++ library for the development of graphical decision models (Druzdzel, 1999). Therefore, in comparison to Weka, it is limited to Bayesian decision models and has no possibility for other data mining algorithms. Since it is the most generic approach and suitable for most applications we decided to use the Bayesian search as the algorithm of choice for GeNIE. In effect, heuristic algorithms such as Tree Augmented Naive Bayes are only recommended for large scale projects (Decision Systems Laboratory, 2016).

3.3 Model data

3.3.1 Data acquisition

The accuracy of the developed model depends highly on the underlying data. For training purposes, a real dataset consisting of patients affected by AKI and those not affected by it was needed. This set was obtained from the MIMIC II database from PhysioNet (Lehman et al., 2011) which contains data from intensive care units (ICU) from hospitals in the United States. We utilized the contained information about disease indications, demographics, lab results (most importantly creatinine value measurements) and comorbidities. AKI is represented by the ICD (International Classification of Diseases) code 584.9. The final step

was generating a comma-separated values file by querying the database tables for preprocessing.

3.3.2 Data preprocessing

Besides demographic information, the MIMIC II database provides information about several risk factors. Furthermore, there are tables for medication and laboratory events which were used for the model as well. In order to train and evaluate them, we decided to choose a cross-validation approach. This enables training and evaluation within the same data set.

For this purpose, we extracted two different data sets. The first one consists of 6000 patients, the second of 9000. Table 1 shows the distribution of patients with and without AKI in the two data sets. They contain information about the demographics of a patient, their comorbidities, the latest creatinine value changes and an indication whether a patient was diagnosed with AKI or not.

Entries	AKI	No AKI
6000	50% (Stage 1, 2 & 3)	50%
9000	33% (Stage 1, 2 & 3)	67%

Table 1: **Distribution of the two data sets**

For use in our experiments, we needed to preprocess the data. This included a discretization of continuous values, as well as the computation of auxiliar values derived from available information. We computed the estimated Gromerular Filtration Rate (eGFR) according to the existing guidelines (NIDDK, 2016), since this rate is an important indication of overall renal function. Next, we calculated the increase of serum creatinine for each patient across multiple measurements and used this value for categorizing the severity of kidney injury according to the AKIN guideline (Mehta et al., 2007). The data thus preprocessed can be fed into the tools and offers the necessary information to enable risk prediction and result validation.

3.4 Model input and output

The prepared data set is the basis for the Bayesian network. This means that the risk factors presented in section 2.1 (lab values, comorbidities and demographics) are the input variables for training and running the model. The resulting output are the probabilities for the presence of AKI as well as the classification stages of RIFLE and AKIN as inferred from the provided input. It enables the user to employ it for decision support

with the same or other data sets. For illustration purposes, a graphical representation for the second iteration model is provided in .

3.5 First iteration model

In the first iteration, the model included the input variables as indicated in the AKI literature, encompassing laboratory values, patient demographics and comorbidities. These random variables are the input nodes. Each node has its own probability as well as possible posterior probabilities stored which define its impact on defining AKI. These probabilities were automatically trained from the data set obtained from the MIMIC II database. The model concentrates on the AKI as well as the two classification guidelines RIFLE and AKIN. For the sake of brevity, we will not provide a graphical representation of the first iteration model. The structure, however, will be clear from analyzing the second iteration model, which already incorporated expert feedback.

3.6 Second iteration model

In the next step, we discussed the model in the first iteration with nephrologists from the Charité hospital in Berlin. The following main insights were gained:

1. RIFLE guideline is not used in practice any more since it is an older classification system. Moreover, AKIN is based on RIFLE and is thus the only classification system needed;

2. There are further influence factors which need to be considered. These factors are weight, urethritis and medication history;

3. The time of comorbidities has an influence on making the correct diagnosis. Diseases that are years ago have less impact than more recent diseases;

4. Physicians normally do not trust any systems but only themselves. A CDSS needs to provide demonstrable value. Additionally, nephrologists mostly do not need such a system since they recognize the symptoms themselves based on their experience. A better use case is the ICU, where the physicians are not kidney experts and are overwhelmed with monitoring data.

Based on the discussion at the Charité, we develop a second model, shown in Figure 4. As per the feedback, we removed the node for RIFLE. Furthermore, we added the nodes for the new risk factors. Finally, we appended more dependencies between the factors and AKI as well as AKIN.
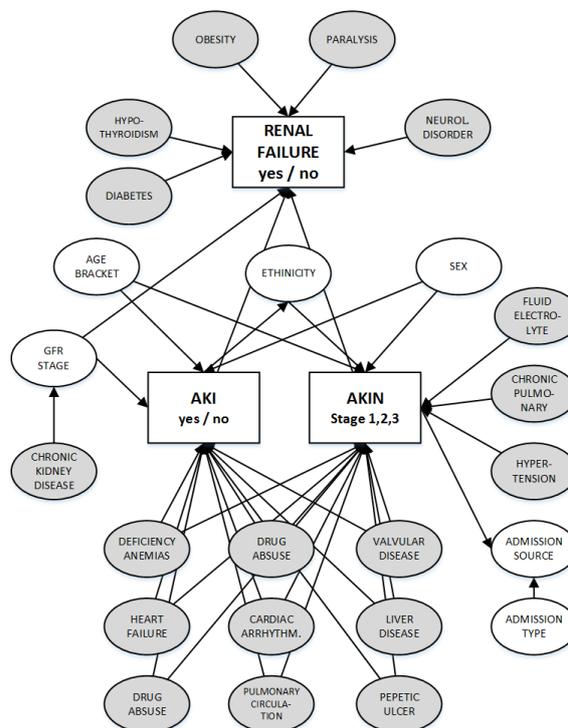

Figure 4: **Second iteration model** Based on the main insights from the expert consultation, this model was developed. The RIFLE guideline was removed. Furthermore, new risk factors and dependencies were added. Comorbidities are shown in gray.

4 Results

This section evaluates how accurate the developed models were. To this extent, the used data set is presented. Moreover, we compare the accuracy values for the two developed models. We show that both improving the model with expert knowledge as well as increasing the data size increases the accuracy of the model. Our best result was an accuracy of 87% for predicting the occurrence of AKI.

4.1 Accuracy of the models

Table 2 shows the obtained results depending on both the utilized tool and the data set used (6000 or 9000). It stands out that the 2nd iteration model consistently performs better than the first one. This demonstrates that expert knowledge is help-

ful in improving model performance in such a specialized scenario as kidney disease. Moreover, it shows the flexibility of Bayesian networks, which allows to integrate such expert knowledge. Furthermore, noticeable discrepancies between the different tools can be observed. Overall, we achieved a top accuracy of 87% when using GeNIe.

	1st iteration		2nd iteration	
Dataset	GeNIe	Weka	GeNIe	Weka
6000	67%	58%	83%	76%
9000	73%	72%	87%	83%

Table 2: **Accuracy of the developed models as per the different tools and datasets**

A more detailed view of the accuracy can be seen by analyzing the receiver operating characteristic (ROC) of the best performing experiment. The ROC curve describes the relation between true positives (TP) and false positives (FP). The perfect result would be a TP value of 100% and a FP value of 0%. Figure 5 shows the curve of the 2nd iteration model (after expert feedback). The curve is based on the larger data set (9000 entries) and refer to computed AKI patients.
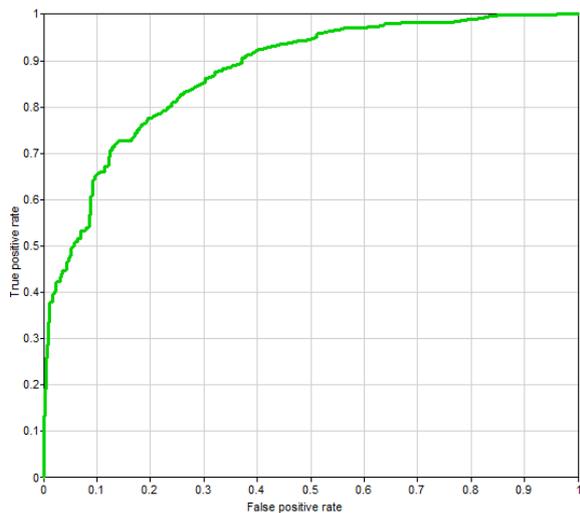

Figure 5: **ROC curve of the 2nd iteration model** Curve shows the performance of the Bayesian classifier for the AKI output variable. Area under the Curve (AUC) = 0.8743

5 Discussion

The results obtained show that for the cohort under analysis increasing data volumes lead to improved accuracy. As such, the higher the volume of data available, the better the results achieved. This was demonstrated by increasing in 50% the data volume (from 6000 to 9000). This finding suggests that more robust prediction models can be developed for the medical domain by tapping into larger databases.

The results also show discrepancies concerning the tools used (Weka and GeNIe). This can be accounted for by differences in algorithm implementation and configuration parameters. This fact underscores the need for comparison not only among algorithms, but also among different tools and configurations, since the details of algorithm implementation can greatly vary. In this paper, while variations were present, the results were largely consistent, except for a much poorer result from Weka when dealing with the smaller dataset.

The comparably small data set due to limited hardware capacity was the biggest drawback of our experiments. Furthermore, the limited scope of this work lead to the decision of concentrate on one algorithm per tool. Indeed, different algorithms show different advantages which we were not able to consider for this paper. As such, a more robust analysis should include a comparison of different algorithms, tools and configuration parameters.

Going one step further, instead of solely increasing the data volume, another promising direction to follow would be to increase data variety, including more relevant data, such as urethritis as a comorbidity and complete disease history. As the experts consulted suggested, this might improve the accuracy as well. Unfortunately, this data could not be obtained from the data source available but it must be included if this this model is to be used in practice.

In the beginning, we showed related works that developed a CDSS for various diseases (Onisko et al., 2004; Nachimuthu and Haug, 2012; Ward et al., 2014). These works employ dynamic Bayesian networks, which consider the time component as the involved random variables change. The approach presented in this work is concerned with a static view. This represents a possible weakness in comparison to other similar works and must be addressed in the future.

Finally, while decision support systems show much promise in improving healthcare delivery, evidence towards their efficacy in clinical settings

is lacking, leading to skepticism among medical professionals. Particularly in the field of Nephrology, a controlled randomized trial (CRT) has been conducted by Wilson et al. testing an early warning alert system for AKI. The CRT which yielded no demonstrable positive outcomes for patients (Wilson et al., 2015). Their algorithm was based on the mere detection of creatinine thresholds and the authors of this study encouraged new trials with more sophisticated algorithms. Such experiences strengthen the need for making CRT a standard procedure for a prospective CDSS. Even though such procedures are costly, if benefits can be factually demonstrated, medical acceptance can be increased.

5.1 Further work

Since the results show that larger data sets tend to deliver more accurate results, the tests should be repeated with more representative data sets that might also contain new input variables like urethritis as a comorbidity and the history of diseases. Furthermore, in a practical application, the machine learning system has to be easily modifiable. One approach is to develop a Clinical Decision Support System specifically for this purpose. Another one is to use the existing tools, but perform more experiments with changing parameters.

For actual use in practice, integration with care delivery workflows and the hospital information system itself is needed. The user interface of such a system has to be clearly structured and free of unnecessary information, so that physicians can readily see the most important facts about a patient. Further, an alert system powered by the clinical decision support system must be implemented, so that care team can be notified timely in case of AKI. Still, it remains to be seen how many physicians would actually use such a system, so acceptance is a possible barrier. Similarly, the expected benefits for the patients must also be measured in terms of outcomes (rate of complication, mortality, prognosis, subjective well-being, etc.) Therefore, a pilot study considering both physician acceptance and patient outcomes should be conducted.

6   Conclusion

We developed a proof of concept for a machine learning model that can be used in a CDSS in the domain of AKI. For this purpose, we trained a Bayesian network on 9000 data entries from ICUs in the US, obtained from the MIMIC database. Using information about demographics, comorbidities and creatinine values, a satisfactory accuracy for predicting the risk of an AKI was obtained. The results show that a further development and improvement of such a model by integration of expert knowledge leads to improved accuracy values. However, such initiatives are frequently met with skepticism by the medical community. Randomized controlled trials are needed to assess the benefits and potential risks for patients and doctors, along with full integration with medical workflows, so that they can be convinced of the potential advantages of such a system.

7   Acknowledgements

References

Emanuela Barbini, Pietro Manzi, and Paolo Barbini. 2013. Bayesian Approach in Medicine and Health Management. In *Current Topics in Public Health.* InTech, may.

Rinaldo Bellomo, Claudio Ronco, John A Kellum, Ravindra L Mehta, and Paul Palevsky. 2004. Acute renal failure - definition, outcome measures, animal models, fluid therapy and information technology needs: the Second International Consensus Conference of the Acute Dialysis Quality Initiative (ADQI) Group. *Critical care (London, England)*, 8(4):R204–12.

Remco Bouckaert. 2008. Bayesian network classifiers in Weka for version 3-5-7. Technical report.

Dinna N Cruz, Zaccaria Ricci, and Claudio Ronco. 2009. Clinical review: RIFLE and AKIN–time for reappraisal. *Critical care (London, England)*, 13(3):211.

Decision Systems Laboratory. 2016. GeNIe Documentation.

Marek J Druzdzel. 1999. SMILE: Structural Modeling, Inference, and Learning Engine and GeNIe: A Development Environment for Graphical Decision-Theoretic Models. *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, pages 342–343.

John Harty. 2014. Prevention and management of acute kidney injury. *The Ulster medical journal*, 83(3):149–57.

D Heckerman, D Geiger, and D Chickering. 1995. Learning Bayesian networks: the combination of knowledge and statistical data. *Machine Learning, 20, 197-243*, 243:197–243.

Michal Horný. 2014. Bayesian networks: A Technical report. Technical Report 5.

Ronald A. Howard and James E. Matheson. 1983. *Readings on the principles and applications of decision analysis*. Sdg Decision Systems, Menlo Park.

John a Kellum, Norbert Lameire, Peter Aspelin, Rashad S Barsoum, Emmanuel a Burdmann, Stuart L Goldstein, Charles a Herzog, Michael Joannidis, Andreas Kribben, Andrew S Levey, Alison M MacLeod, Ravindra L Mehta, Patrick T Murray, Saraladevi Naicker, Steven M Opal, Franz Schaefer, Miet Schetz, and Shigehiko Uchino. 2012. KDIGO Clinical Practice Guideline for Acute Kidney Injury. *Kidney international supplements*, 2(1):1–138.

Ewa Król, Bolesław Rutkowski, Piotr Czarniak, Ewa Kraszewska, Sławomir Lizakowski, Radosław Szubert, Stanisław Czekalski, Władysław Sułowicz, and Andrzej Wicek. 2009. Early detection of chronic kidney disease: Results of the PolNef study. *American Journal of Nephrology*, 29(3):264–273.

Yugal Kumar and G. Sahoo. 2012. Analysis of Bayes , Neural Network and Tree Classifier of Classification Technique in Data Mining using WEKA. *Computer Science & Information Technology (CS & IT)*, pages 359–369.

Matthieu Legrand, Romain Pirracchio, Anne Rosa, Maya L Petersen, Mark Van der Laan, Jean-Noël Fabiani, Marie-paule Fernandez-gerlinger, Isabelle Podglajen, Denis Safran, Bernard Cholley, and Jean-Luc Mainardi. 2013. Incidence, risk factors and prediction of post-operative acute kidney injury following cardiac surgery for active infective endocarditis: an observational study. *Critical care (London, England)*, 17(5):R220, jan.

Li-wei Lehman, George Moody, Thomas Heldt, and Tin H Kyaw. 2011. Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC- II): A public-access intensive care unit database. *Critical Care*, 39(February 2010):952–960.

José António Lopes and Sofia Jorge. 2013. The RIFLE and AKIN classifications for acute kidney injury: A critical and comprehensive review. *Clinical Kidney Journal*, 6(1):8–14.

Ravindra L. Mehta, John A. Kellum, Sudhir V. Shah, Bruce A. Molitoris, Claudio Ronco, David G. Warnock, Adeera Levin, and $author firstName $author.lastName. 2007. Acute Kidney Injury Network: report of an initiative to improve outcomes in acute kidney injury. *Critical Care*, 11(2):R31.

Senthil K Nachimuthu and Peter J Haug. 2012. Early detection of sepsis in the emergency department using Dynamic Bayesian Networks. *AMIA ... Annual Symposium proceedings / AMIA Symposium. AMIA Symposium*, 2012:653–62, jan.

NIDDK. 2016. Estimating Glomerular Filtration Rate (GFR).

A Onisko, MJ Druzdzel, and RM Austin. 2004. Application of Dynamic Bayesian Networks to Cervical Cancer Screening. *Intelligent Information Systems*, pages 1–10.

Zaccaria Ricci, Dinna N Cruz, and Claudio Ronco. 2011. Classification and staging of acute kidney injury: beyond the RIFLE and AKIN criteria. *Nature reviews. Nephrology*, 7(4):201–8, apr.

Statistisches Bundesamt. 2014. Todesursachen in Deutschland.

J Stewart, G Findlay, N Smith, K Kelly, and M Mason. 2009. Adding insult to injury. *National Confidential Enquiry into Patient Outcome and Death*, s10-I(1):4.

The University of Waikato. 2008. Search algorithms.

Logan Ward, Mads L Mogensen, Mical Paul, Leonard Leibovici, and Steen Andreassen. 2014. A Bayesian Approach to Model-Development : Design of Continuous Distributions for Infection Variables. In *19th IFAC World Congress*, pages 5653–5658.

F Perry Wilson, Michael Shashaty, Jeffrey Testani, Iram Aqeel, Yuliya Borovskiy, Susan S Ellenberg, Harold I Feldman, Hilda Fernandez, Yevgeniy Gitelman, Jennie Lin, Dan Negoianu, Chirag R Parikh, Peter P Reese, Richard Urbani, and Barry Fuchs. 2015. Automated, electronic alerts for acute kidney injury: a single-blind, parallel-group, randomised controlled trial. *The Lancet*, 385(9981):1966–1974.