

# Subpopulation process mining in healthcare

Simona Filipović  
University of Twente  
P.O. Box 217, 7500AE Enschede  
The Netherlands  
s.filipovic@student.utwente.nl

## ABSTRACT

In clinical pathways, there are thought to be differences between the treatment of different patient subpopulations. This paper provides a method for comparing clinical pathways of different patient subpopulations. To perform and validate the method three diseases were chosen, diabetes type II, chronic kidney disease and urinary tract infection. Within these diseases a number of different subpopulations were chosen from the MIMIC-III v1.4 data set to be compared against each other. Analysis of data shows statistically significant differences in clinical pathways in the form of graph per subpopulation. The results indicate that it is possible to apply process comparison within process mining on medical data and that the resulting models are sound within the medical domain.

## Keywords

process mining, healthcare, subpopulation, clinical pathways

## 1. INTRODUCTION

Healthcare Information Systems (HIS) have hundreds of tables with patient-related event data. [1] This data can be used to potentially improve healthcare procedures, thus, in turn, improve patient care and treatment, by implementing process mining techniques. Many markers define a person's clinical pathway in healthcare. Through these markers, different subpopulations of patients can be identified. One of those markers could be gender which creates the division of patients into a male and a female subpopulation. Other markers could also be age, religion, ethnicity, insurance, vitals etc. It is important to find these markers and establish what the optimal paths for the subpopulations are in order to provide the best care possible. In this paper, the focus lies on providing a method for identifying similarities and differences of different subpopulations for one specific disease.

## 2. BACKGROUND

### 2.1 Process mining

Process Mining focuses on extracting knowledge from data generated and stored in (corporate) information systems

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

33<sup>rd</sup> Twente Student Conference on IT July 3<sup>rd</sup>, 2020, Enschede, The Netherlands.

Copyright 2020, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science.

in order to analyze executed processes.[2] There are three phases to process mining: process discovery, conformance checking and enhancement. Process discovery represents the model extraction from logs, conformance checking is the comparison of the retrieved model to the event log and enhancement is the improvement of the process through the knowledge acquired through process mining. The event log for process mining must include:

- Case ID - an identifier to determine different executions of the same process.
- Activity - steps that are performed during the process.
- Timestamp - exact moments when the activity steps took place.

### 2.2 Tools

The tool used during the research project was ProM, an extensible framework that supports a wide variety of process mining techniques in the form of plug-ins[3]. The plugin used for the project was process comparator. The plugin is able to detect relevant differences undetected by previous approaches while it avoids detecting insignificant differences.[4]

In order to extract the data the Query Builder provided by the MIT Laboratory for Computational Physiology was initially used. It provided the ability to export the results of necessary queries for further processing. While working with the Query Builder the limitations of its data extraction capabilities were observed thus the extraction of the data was moved to the Big Query services provided by the Google Cloud Platform. When it came to data processing additional scripts using python pandas library were written.

### 2.3 Data

The data used in this research project is from the MIMIC-III Clinical Database. It is a freely-available database comprising de-identified health-related data associated with over forty thousand patients who stayed in critical care units of the Beth Israel Deaconess Medical Center between 2001 and 2012.[5] It encompasses a diverse and very large population of ICU patients making it a representative database. MIMIC-III consists of 26 different relational tables.

## 3. RESEARCH QUESTION

The research questions posed in this paper are:

1. How to identify suitable subpopulations of a disease in a data set?
2. How to determine differences in clinical pathways of suitable subpopulations per disease?

## 4. RELATED WORK

Up to 2016, there have been 74 different papers published in the field of process mining in healthcare. Aspects that these papers explored were: process and data types, methodologies, process mining perspectives, algorithms and so on. The most used techniques according to the review were the heuristic miner, the fuzzy miner and trace clustering. A number of case studies also took place of which the majority were in the oncology, surgery and cardiology field.[2]

Furthermore, in 2018, 55 articles were covered and showed that 29% of papers focused on the comparison of processes and 12% focused on process mining on clinical pathways. [6] Process mining has also been applied to stroke care with two data sets, one being the clinical course of the patient and the other the pre-hospital behaviour data of the stroke patients to identify clinical pathways and bottlenecks.[7]

Other research within process mining done on the same data set showed that it is possible to mine complex medical processes with current algorithms to discover and analyse process models.[8]

In 2019 research comparing processes for different patient populations in breast cancer care was conducted. The populations were divided based on age, BIRAD score and whether the patients were sent by a general practitioner or national breast cancer screening program. The research showed that average fitness and precision of cross-log conformance checks provide good indications of process similarity. [9]

## 5. METHODOLOGY

For this research, the methodology was partially based on the methodology proposed by article PM<sup>2</sup>: a Process Mining Project Methodology [10].

### 5.1 Stage 1: Planning

The first step is meant for setting the goals and the questions of the research. Preferably, in this stage, you have domain experts that are willing to help you understand the data so that the goals and questions are achievable.

### 5.2 Stage 2: Data extraction

Stage 2 represents data extraction from the database. Whether questions set in step 1 can be answered will depend on the availability and the quality of the data. In this stage, it is also important to think about which activities should be included in the final model. This data will be used to build the event log necessary for stage 4.

### 5.3 Stage 3: Data processing

During stage 3 the extracted data is processed. This involves formatting data into a usable event log. This includes setting the case ID to the patient's ID and finding activities and timestamps for each of the patients in the subpopulations. Further tidying up of data includes deleting duplicate rows and removing incomplete data. Additionally, in this stage, if necessary, additional techniques for dealing with wrongly input data can be applied.

### 5.4 Stage 4: Process mining

In stage 4 process mining is applied to answer the posed research questions. During stage 4 ProM plugin Process Comparator is applied to determine statistically significant differences in specific procedures of clinical pathways. For the results, the "hint" function is used to calculate a similarity score between subpopulations. The similarity score is calculated based on the percentage of elements that present a statistically significant difference.

## 5.5 Stage 5: Evaluation and Summarizing results

In the final stage, the results are summarized and reported. These reports are used to interpret the findings with domain experts. Following the meeting with the domain experts a reiteration of summarisation of the resulting models and key findings together is crucial as the answers to the posed research questions are then obtained.

The article methodology also includes a step of process improvement. As this research project is not done in a collaboration with a specific hospital where the changes will be implemented that particular step will not take place in this research.

## 6. EXPERIMENTAL SET-UP

### 6.1 Stage 1: Planning

The planning stage of the research project manifested itself through the work on the research proposal. Before the research questions and the goals could be set a certain familiarization with the data was needed. This was done through the tools offered by the MIT Laboratory for Computational Physiology, such as QueryBuilder and the schema of the database itself, as well as, thorough readings of the offered documentation. Once an idea was formed on how to handle the data, the task of setting the goals and questions was completed.

### 6.2 Stage 2: Data extraction

During the process of data extraction multiple diseases needed to be selected. The number of suitable diseases was set to three in order to validate the research method. The criteria of what is a suitable subpopulation was based on whether it is representative, namely that any part of subpopulation had enough patients in order for it to be possible to mine for the clinical pathways. The best approach to identify suitable subpopulations was trial and error with preemptive research in most common diseases and, primarily, discussions with medical experts. Consultation with domain experts vastly helped when it came to suggestions where to start searching for diseases with representative subpopulations, as well as what those subpopulations might be. The diseases that were found to suitable within the MIMIC-III v1.4 data set were diabetes type II, chronic kidney disease and urinary tract infection.

In order to extract the data SQL queries were adopted. For every disease at least one ICD 9 code<sup>1</sup> had to be selected in the query in order to select patients for the necessary disease. Different patient subpopulations required different specified conditions. Certain markers, such as gender, admission type and length of stay in the ICU were existing values in the database and extracting these subpopulations did not need anything more than simple WHERE clauses to separate the patient subpopulations. Age of patients as a value was not stored in the database itself, however, dates of birth and dates of admission were. The difference between these values was used to get the patient age at the time of admission thus making the mining of the age subpopulation possible. Additionally, values of creatinine/bacteria levels were accessed through identifiers for laboratory events. The separation of those subpopulations was done with the use of flags. The values for creatinine/bacteria levels, apart from being exact, also had an

<sup>1</sup>The International Classification of Diseases Clinical Modification, 9th Revision is a list of codes intended for the classification of diseases. The numerical format of the diagnosis codes usually ranges from three to five digits that are assigned to a unique category.

accompanying flag which was set to abnormal if the results were anomalous. With the use of these flags subpopulations normal/abnormal were made for all laboratory events. Lastly, when it came to BMI, patient weight was stored as a value in the inpuvents table. Many attempts of calculating the BMI were made by using the average height of men and women in the state of Massachusetts. The averages of this state were chosen as the hospital was located there. With further immersion with the data set, the code for the event of getting a patient's height was found in the procedures table. All of the procedures selected in the final models had unique procedure identifiers in the same table as the identifier for the patient height that needed to be selected. In order for it to be possible to calculate the BMI nested queries were applied as a solution. In order to calculate the body mass index the following formula was used  $BMI = \frac{kg}{m^2}$ , where kg is a patient's weight in kilograms and  $m^2$  is their height in metres, squared. Values of BMI over 30 represented the obese subpopulation.

The extracted data comprised of patient identifiers, clinical procedures done on the respectful patient as well as starting time and ending time of said procedures.

### 6.3 Stage 3: Data processing

Once the data was extracted further polishing was necessary. As the extracted data was already somewhat in an appropriate form, small changes needed to be made in order to have files ready for process mining. The taken measures included deletion of duplicate and incomplete rows. Certain data inaccuracies were encountered during this stage. Most mistakes that were noticed were obvious such as patient weights being 1 kilogram. This was kept in mind and python scripts that removed patients if they appeared in both of the subpopulations were applied. This was done for the case of BMI, but as well as with volatile values such as creatinine/bacteria levels. It was noticed that during the stay these values were oscillating between normal and abnormal for certain patients. In order to have models that would represent only one of the normal/abnormal subpopulation deletion of patients was done if they appeared in both since that would result in more sound models.

### 6.4 Stage 4: Process mining

Once the files were ready it was time to start process mining. ProM plugin process comparator was used as the objective of the plugin is to find statistically significant differences between variants of the same process. One of the features was to filter the elements below a certain threshold. This was used in order to remove outliers in medical procedures. This threshold was set to 5% on most of the models with an exception which included the BMI medical prescriptions for diabetes type II. The frequency threshold, in this case, was set to 20% as the resulting model with 5% was illegible and in-comprehensible. Another feature of the plugin was the alpha significance level. All the resulting models applied a 5% alpha significance level. Furthermore, the hint feature was used to obtain the percentage of statistically significant differences between all of the subpopulations within a specific disease. The metric used to compare the processes was set to occurrence frequency, i.e. the number of times the procedures occurred within the selected data files. Lastly, once all settings were in place the models were extracted.

## 7. RESULTS

### 7.1 Disease 1 - Diabetes type II

In order to identify patients with diabetes type II multiple ICD 9 codes were used. The code 25000 specified "Diabetes mellitus without mention of complication, type II or unspecified type, not stated as uncontrolled" and code 25002 specified "Diabetes mellitus without mention of complication, type II or unspecified type, uncontrolled". In total, the number of patients with above mentioned ICD 9 codes was 7654.

The subpopulations chosen for Diabetes type II were:

1. Age
2. Gender
  - Male vs female
  - Male vs female within the emergency subpopulation
3. Admission type
4. Length of stay in ICU
5. BMI
  - Procedures
  - Medication

**Table 1. Diabetes type II: statistically significant differences in clinical pathways per subpopulation**

Subpopulation	Results
Age	65+ vs 45-64: 4.40%
	65+ vs 18-44: 3.53%
	45-64 vs 18-44: 3.13%
Gender	Male vs female: 24.18%
	Male vs female within the emergency subpopulation: 18.75%
BMI	BMI over 30 VS BMI under 30 in procedures: 4.55%
	BMI over 30 VS BMI under 30 in medical prescriptions: 12.19%
Stay in ICU	ICU stay over 3.4 days vs ICU stay under 3.4 days: 75.28%
Admission type	Emergency vs Elective: 51.69%
	Urgent vs Elective: 13.89%
	Urgent vs Emergency: 1.23%

Obesity has a strong relationship with diabetes type II and insulin resistance [11]. Subsequently, differences larger than 4.55%, as seen in Table 1 were expected in the subpopulation models when it came to BMI. In order to explain such a small difference, a closer look was taken into the treatment of the BMI subpopulation with respect to the medication administered to patients. In the models with medication prescription statistically significant difference went up to 12.19%, as well as insights why the procedures had a difference of 4.55%. Heparin, an anti-coagulant medication and furosemide, a medication used to treat fluid build-up due to heart failure, liver scarring, or kidney disease are statistically more administered to the obese subpopulation implying that obese people also suffer more from heart issues than the non-obese subpopulation. Due to this fact, any intensive procedures on the obese subpopulation would not be taken as they are seen as a high risk for the obese subpopulation. Models used to come to these findings are visible in Figure 1 and 2.

Furthermore, the difference in clinical pathways for the gender subpopulation was 24.18%, which is higher than expected. To further look into this a gender separation was

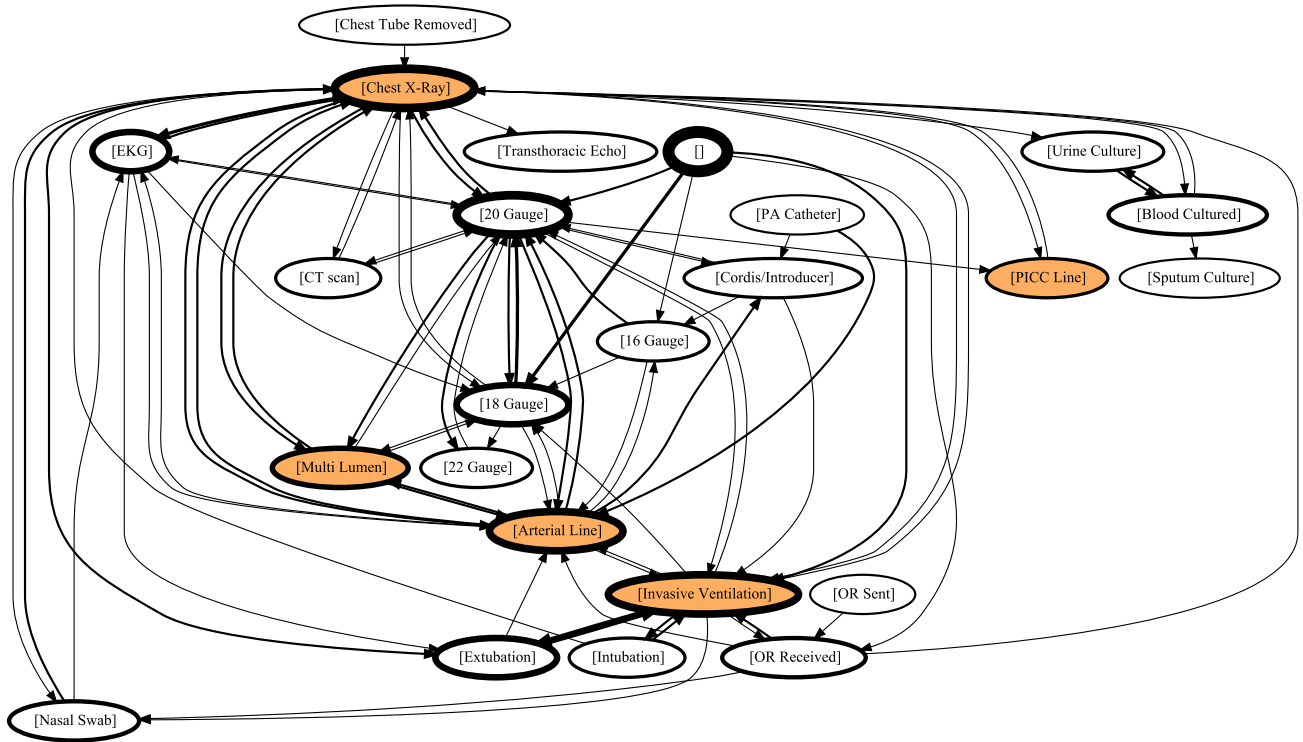


Figure 1. BMI procedure comparison: Occurrence frequency comparison. Coloured nodes and transitions represent statistically significant differences between the two subpopulations. Blue colours represent a higher occurrence in the BMI under 30 subpopulation. Red colours represent a higher occurrence in the BMI over 30 subpopulation. There is no occurrence of blue states which implies that there are no procedures that are more often executed on the subpopulation of BMI under 30 than on the subpopulation of BMI over 30.

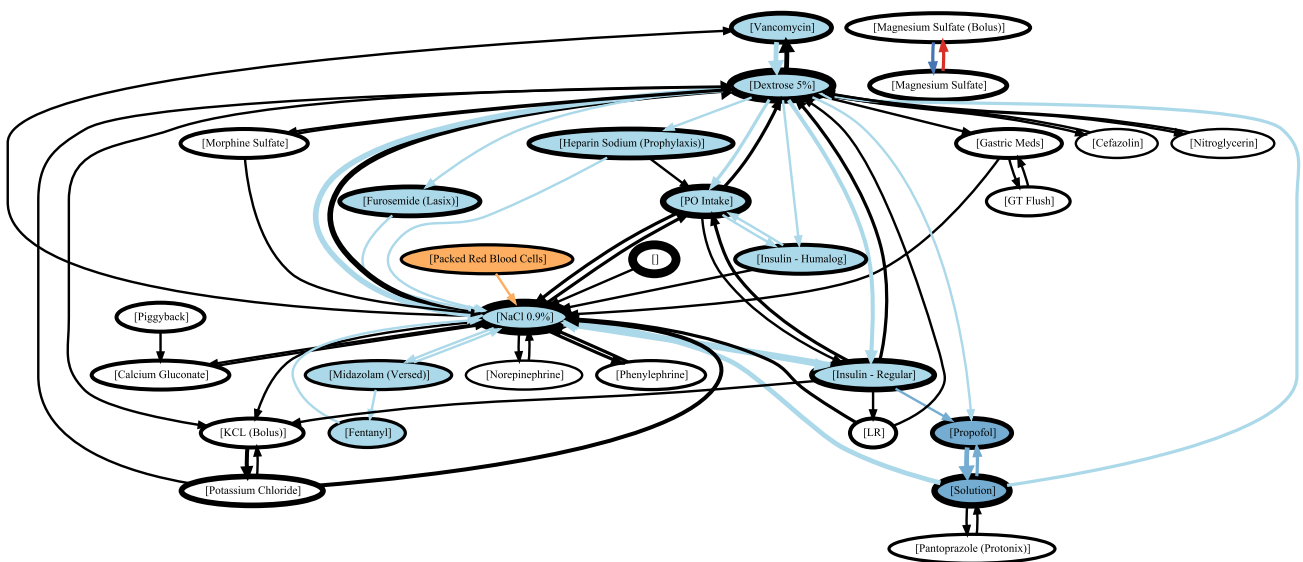


Figure 2. BMI medication prescription comparison: Occurrence frequency comparison. Coloured nodes and transitions represent statistically significant differences between the two subpopulations. Blue colours represent a higher occurrence in the BMI over 30 subpopulation. Red colours represent a higher occurrence in the BMI under 30 subpopulation.

created within the emergency subpopulation, as the sample size of that particular subpopulation was large enough. This brought the difference down to 18.75%.

## 7.2 Disease 2 - Chronic kidney disease

In order to identify patients with a chronic kidney disease multiple ICD 9 codes were used. The codes 5851, 5852, 5853, 5854, 5855, 5856 specified "Chronic kidney disease, Stage I to VI" and code 5859 was used for the "Chronic kidney disease, unspecified". In total, the number of patients with above mentioned ICD 9 codes was 4689. Table 2 summarizes findings of the created models.

The subpopulations chosen for Chronic kidney disease were:

1. Age
2. Gender
3. Admission type
4. Length of stay in ICU
5. BMI
  - Procedures
  - Medication
6. Creatinine levels

**Table 2. Chronic kidney disease: statistically significant differences in clinical pathways per subpopulation**

Subpopulation	Results
Age	65+ vs 45-64: 15.89%
	65+ vs 18-44: 4.29%
	45-64 vs 18-44: 5.13%
Gender	Male vs female: 13.92%
BMI	BMI over 30 vs BMI under 30 in procedures: 0.00%
	BMI over 30 vs BMI under 30 in medical prescriptions: 5.47%
Stay in ICU	ICU stay over 3.4 days vs ICU stay under 3.4 days: 78.87%
Admission type	Emergency vs Elective: 41.67%
	Urgent vs Elective: 13.86%
	Urgent vs Emergency: 4.00%
Creatinine levels	Normal vs abnormal: 4.11%

When it came to creatinine levels it is seen that the subpopulation with high levels of creatinine was more often subjected to dialysis due to worsening kidney function, visible in Figure 3.

## 7.3 Disease 3 - Urinary tract infection

In order to identify patients with the urinary tract infection an ICD 9 code was used. The code 5990 specified "Urinary tract infection, site not specified". The number of patients with this ICD 9 code was 5779. Table 3 summarizes findings of the created models.

The subpopulations chosen for Urinary tract infection were:

1. Age
2. Gender
3. Admission type
4. Length of stay in ICU
5. Bacteria levels in urine

**Table 3. Urinary tract infection: statistically significant differences in clinical pathways per subpopulation**

Subpopulation	Results
Age	65+ vs 45-64: 15.29%
	65+ vs 18-44: 3.85%
	45-64 vs 18-44: 2.94%
Gender	Male vs female: 11.63%
Stay in ICU	ICU stay over 3.4 days vs ICU stay under 3.4 days: 85.56%
Admission type	Emergency vs Elective: 25.58%
	Urgent vs Elective: 6.20%
	Urgent vs Emergency: 4.55%
Bacteria levels in urine	Normal vs abnormal: 22.54%

When it came to the analysis of the bacterial levels in urine models, it was noticed that overall, more procedures are done on patients that had normal levels of bacteria in urine. This is explained by the fact that most urinary tract infections are caused by a type of bacteria. Patients that have high levels of bacteria in their urine were more easily diagnosed whilst the patients that had normal levels of bacteria had a higher occurrence of blood culture, urine culture, different types of scans in order to try to diagnose those patients and find the cause of the infection. This is visible in Figure 4.

## 8. DISCUSSION

### 8.1 Data selection

In the first iteration of the project information such as the admission, discharge, ICU transfer and emergency room registration were taken into account. These type of procedures were omitted in the final models. Firstly, the reason this was done is because this type of data represented more of an administrative practice rather than a clinical pathway of a patient. Secondly, every patient was admitted and discharged from the hospital and thus would not be useful when identifying differences between subpopulations. Furthermore, ICU stays and emergency room registration provided to be better for the separation of the subpopulations than as processes in the clinical pathways of patients.

### 8.2 Medical analysis

Medical experts were consulted in order to gain domain knowledge and explain the resulting models. Overall, certain patterns emerged for some of the subpopulations in all of the tested diseases.

Firstly, the stay in the ICU subpopulation was divided into the subpopulation where the stay in the ICU exceeded 3.4 days and was lower than 3.4 days. This number was picked as research has shown that 3.4 days is the average time spent in the ICU. [12] Models for the stay in the ICU resulted in splitting the subpopulation into more severe and less severe cases in the subpopulation. Subsequently, the differences were large and all of the procedures were statistically significant in the subpopulation where the stay was longer than 3.4 days.

Secondly, in the admission type subpopulations, it was visible that the difference between the elective and urgent/emergency resulted in models that showed that the elective patients were sent for procedures such as surgery, while with the patients marked as urgent/emergency most




Figure 3. Creatinine levels comparison: Occurrence frequency comparison. Coloured nodes and transitions represent statistically significant differences between the two subpopulations. Blue colours represent a higher occurrence in the normal levels of creatinine subpopulation. Red colours represent a higher occurrence in the abnormal levels of creatinine subpopulation.

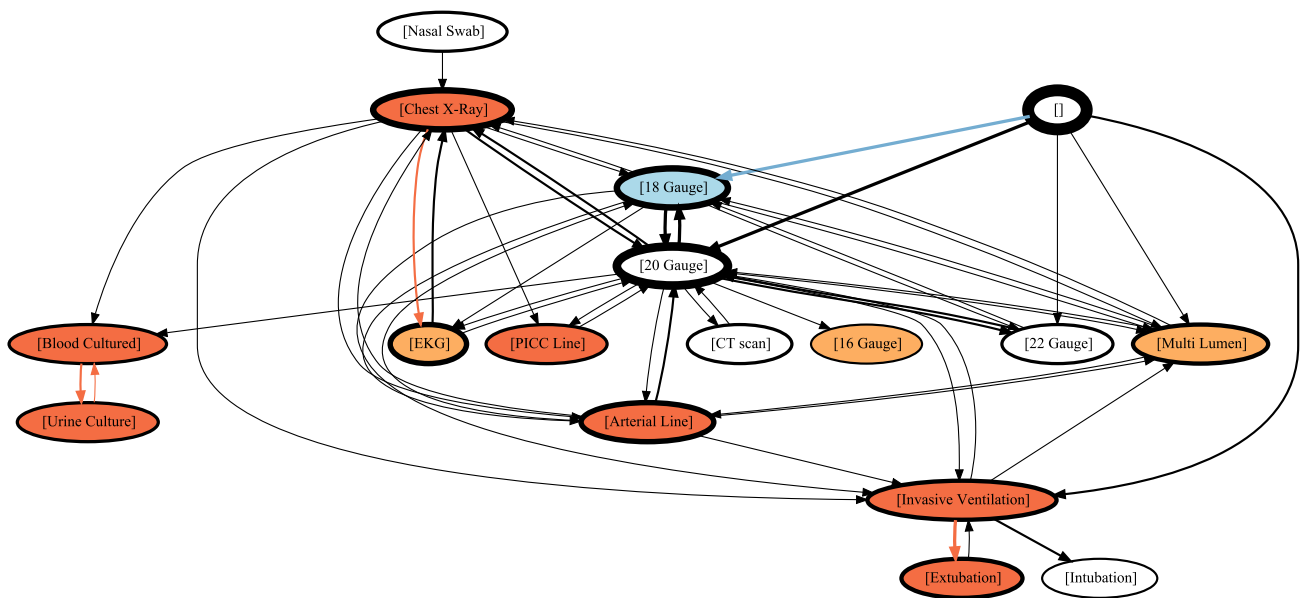


Figure 4. Bacteria level comparison: Occurrence frequency comparison. Coloured nodes and transitions represent statistically significant differences between the two subpopulations. Blue colours represent a higher occurrence in bacteria levels: MANY AND MODERATE subpopulation. Red colours represent a higher occurrence in the bacteria levels: NOT many or moderate subpopulation.

of the procedures were of the test type, such as blood culture, urine culture, CT scan etc. This is explained by the fact that the elective patients were sent in by their general practitioners and came to the hospital with a clear problem, whilst the emergency/urgent patients came in not knowing what the medical issues were and thus a plethora of tests had to be done to diagnostically determine the issues of the emergency/urgent patients at hand.

Furthermore, when it came to the age subpopulation it was visible that more procedures were being done on the subpopulation younger than 65. This is explained by the fact that any intensive procedures are riskier for the older subpopulation than the younger counterpart and that certain procedures are not possible when a certain age is reached.

Lastly, in the gender comparison, a difference that was always identified was the gauge procedure with a number specification. It was observed that in the female subpopulation the gauge number was higher, as that is an indication of a smaller diameter of hypodermic needles. This finding can also be applied to other models in order to infer the possible gender majority in other subpopulations, such as age, admission type etc.

### 8.3 Graph interpretation

The nodes that remained white represent activities that appear in both of the subpopulations and represent no statistically significant differences. The thickness of the transitions represents the trace frequency of certain transitions, the thicker the transition line the higher the trace of that transition is in the data. If the transitions remain black it suggests that there is no statistically significant difference between the two subpopulations. On the other hand, if they are coloured that means that there is a statistically significant difference. Different shades of red and blue represent the statistical significance, darker colours indicate higher statistical significance.




Figure 5. Colour scale in process comparator

## 9. LIMITATIONS AND FUTURE WORK

During the research hardware limitations in terms of memory availability were reached. This happened when trying to expand the model seen in Figure 2 with the amount of the prescription medication the patients were receiving.

Further work on this data set could try to apply data mining techniques on the notes written by caregivers about patients in order to better explain the models, as well as provide a better understanding of patients with volatile markers. Additional analysis of notes with domain experts could lead to patients with volatile markers not to be removed from the models, it could be understood which procedures would belong to which subpopulation. Furthermore, depending on the disease and subpopulations that are selected those patient notes might be essential. This could be particularly seen in certain procedures such as CT scans or MRI scans where result values indicate only "done". Results of these procedures are given in the notes. If this type of research were to be done on neuro-

logical issues, where results of such procedure are highly important an additional in-depth look in the data would be necessary.

## 10. CONCLUSION

To summarize, this research aimed to identify suitable patient subpopulation, as well as to determine the differences in clinical pathways between the suitable patient subpopulations.

In order to discover applicable, i.e. representative, subpopulations the method of trial and error preceded by consultation with medical experts showed effective and convenient. In line with that method, suitable subpopulations within multiple diseases were found in the MIMIC-III v1.4 data set. The diseases that comprised of suitable subpopulations included diabetes type II, chronic kidney disease and urinary tract infection.

As to uncover the differences in each of the subpopulations' clinical pathways per disease the methodology PM<sup>2</sup> was followed. With the use of the ProM plugin process comparator statistically significant differences were discovered. Results showed that the proposed method can be used to identify the differences and similarities of clinical pathways and that the resulted models are sound within the medical domain.

Results have shown that there are some omnipresent findings within certain subpopulations across the researched diseases.

Firstly, the admission type subpopulations findings showed that the subpopulations differ due to a contrasting medical diagnosis procedure. Urgent/emergency patients needed more procedures focusing on determining the condition. This difference, the focus on the medical diagnostic, is also seen in a disease-specific subpopulation, bacteria levels in the urinary tract infection. Given urinary tract infections are mostly caused by abnormal bacteria levels, normal levels would then indicate a different underlying cause for the disease. Therefore, a more extensive diagnostics check is necessary.

Furthermore, the age subpopulation differed due to the heightened risk to perform invasive medical procedures on older patients. This situation is repeated in the disease-specific BMI subpopulation in diabetes type II, as the more obese subpopulation also suffers from a higher risk of cardiovascular issues. Therefore, a medically invasive procedure represents a precariousness.

Lastly, the length of stay in the ICU subpopulations differed based on the severity of the patient condition. This difference, the gravity of the patient's state, is also seen in a disease-specific subpopulation, creatinine levels in chronic kidney disease, where the patients with abnormal creatinine levels had a more severe case.

## 11. REFERENCES

- [1] Ronny S Mans, Wil M P Van Der Aalst, and Rob J B Vanwersch. *Process Mining in Healthcare Evaluating and Exploiting Operational Healthcare Processes*.
- [2] Eric Rojas, Jorge Munoz-Gama, Marcos Sepúlveda, and Daniel Capurro. Process mining in healthcare: A literature review. *Journal of Biomedical Informatics*, 61:224–236, 2016.
- [3] start | ProM Tools.
- [4] Alfredo Bolt, Massimiliano de Leoni, and Wil M. P. van der Aalst. A Visual Approach to Spot

- Statistically-Significant Differences in Event Logs Based on Process Metrics. pages 151–166. 2016.
- [5] MIT Laboratory For Computational Physiology. The MIMIC III Clinical Database. pages 1–14, 2015.
- [6] Edgar Batista and Agusti Solanas. Process mining in healthcare: A systematic review. In *2018 9th International Conference on Information, Intelligence, Systems and Applications, IISA 2018*. Institute of Electrical and Electronics Engineers Inc., feb 2019.
- [7] Ronny Mans, Helen Schonenberg, Giorgio Leonardi, Silvia Panzarasa, Anna Cavallini, Silvana Quaglini, and Wil Van Der Aalst. Process Mining Techniques: an Application to Stroke Care. Technical report.
- [8] Angelina Prima Kurniati, Geoff Hall, David Hogg, and Owen Johnson. Process mining in oncology using the MIMIC-III dataset. *Journal of Physics: Conference Series*, 971(1), 2018.
- [9] Francesca Marazza, Faiza Allah Bukhsh, Onno Vijlbrief, Jeroen Geerdink, Shreyasi Pathak, Maurice van Keulen, and Christin Seifert. Comparing Process Models for Patient Populations: Application in Breast Cancer Care. *Lecture Notes in Business Information Processing*, 362 LNBIP:496–507, 2019.
- [10] Maikel L. Van Eck, Xixi Lu, Sander J.J. Leemans, and Wil M.P. Van Der Aalst. PM2: A process mining project methodology. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9097:297–313, 2015.
- [11] Abdullah S. Al-Goblan, Mohammed A. Al-Alfi, and Muhammad Z. Khan. Mechanism linking diabetes mellitus and obesity. *Diabetes, Metabolic Syndrome and Obesity: Targets and Therapy*, 7:587–591, dec 2014.
- [12] Vivek K. Moitra, Carmen Guerra, Walter T. Linde-Zwirble, and Hannah Wunsch. Relationship between ICU Length of Stay and Long-Term Mortality for Elderly ICU Survivors. *Critical Care Medicine*, 44(4):655–662, apr 2016.