Aalto University
School of Science

Master’s programme in Mathematics and Operations Research

Process Mining Patient Pathways in Pediatric Emergency Departments at HUS

Tuomo Antikainen

Master’s Thesis
2025

Aalto University
School of Science

Author Tuomo Antikainen

Title Process Mining Patient Pathways in Pediatric Emergency Departments at HUS

Degree programme Mathematics and Operations Research

Major Systems and Operations Research

Supervisor Prof. Lauri Saarinen

Advisor Doc. Miika Koskinen

Collaborative partner HUS Helsinki University Hospital

Date 13 February 2025 **Number of pages** 70 **Language** English

Abstract

Healthcare worldwide is grappling with the challenge of overcrowding, that is caused by the demand for services exceeding the available supply. Various strategies have been implemented to address this issue, with a central objective being the optimization of patient pathways within healthcare systems. This involves ensuring timely and appropriately scaled interventions, thereby enhancing the efficiency of resource utilization. However, this remains challenging due to the dynamic and complex nature of healthcare processes. Alongside widely adopted methods such as Lean and Evidence-Based Management, Process Mining has emerged as a valuable tool to gain novel insights into healthcare processes, providing a better understanding of the efficiency of patient pathways.

This thesis demonstrates how process mining can be utilized with large scale patient data. Prior to this study, process mining had primarily been used with small, well-defined cohorts of patients. In this study, process mining is utilized to analyze patient pathways in the pediatric and adolescent emergency department at HUS Helsinki University Hospital, providing insight into resource allocation, variations in care pathways, and failure demand within the ESI triage system. The process mining analysis is conducted based on ESI triage categories to examine the variations between them. In addition, the thesis explores how complex large-scale healthcare data is feasible to preprocess, which mining techniques are suitable for such data, and how the results can be analyzed from a healthcare management perspective.

The thesis demonstrates that process mining is a valuable technique even with large scale healthcare data. The results show that most pediatric and adolescent patients at the HUS emergency department require minimal resources beyond human resources and follow similar care pathways. An exception to this are the most acute patients in the ESI 1 category, who require substantial resources and follow individualized care pathways. However, the similarities observed across ESI categories suggest that the ESI classification system may not be functioning as intended.

Keywords Process mining, Pediatric ED, Healthcare management

Aalto-yliopisto
Perustieteiden
korkeakoulu

Tekijä Tuomo Antikainen

Työn nimi HUS lasten päivystyksen potilaspolkujen analysointi prosessilouhintamenetelmillä

Koulutusohjelma Matematiikka ja operaatiotutkimus

Pääaine Systeemi- ja operaatiotutkimus

Työn valvoja Prof. Lauri Saarinen

Työn ohjaaja Dos. Miika Koskinen

Yhteistyötaho HUS Helsingin yliopistollinen sairaala

Päivämäärä 13. helmikuuta 2025 **Sivumäärä** 70 **Kieli** englanti

Tiivistelmä

Maailmanlaajuisesti terveydenhuoltojärjestelmien yksi suurimmista haasteista on ylikuormitus, joka johtuu terveydenhuollon palveluiden tarpeiden kasvusta. Ongelman ratkaisemiseksi on kehitetty erilaisia strategioita. Näistä yksi keskeisimmistä on potilaspolkujen tehostaminen. Tällä tarkoitetaan oikea-aikaisia ja järkevästi mitoitettuja hoitointerventioita, joiden avulla pyritään resurssien tehokkaampaan kohdentamiseen. Haasteena on kuitenkin terveydenhuollon prosessien dynaaminen luonne. Laajasti käytettyjen menetelmien, kuten Leanin ja näyttöön perustuvan johtamisen ohella, prosessilouhinnasta on muodostunut kiinnostava työkalu terveydenhuollon prosessien mallintamiseen ja analysointiin. Se tarjoaa paremman käsityksen potilaspolkujen tehokkuudesta.

Tässä diplomityössä selvitetään kuinka prosessilouhintaa voidaan hyödyntää suurten potilasdatajoukkojen analysointiin. Työssä käytetään prosessilouhintaa HUS:n lasten ja nuorten päivystyksen potilaspolkujen mallintamiseen tuomalla esiin tietoa resurssien jakautumisesta, hoitopolkujen vaihteluista ja ESI triage -järjestelmän tehokkuudesta. Potilasdata on luokiteltu ESI triage -kiireellisyysarvojen mukaan, mikä mahdollistaa ryhmien välisten erojen arvioinnin. Lisäksi tässä työssä tutkitaan, millaista esikäsittelyä pirstaloitunut terveydenhuollon tuottama data vaatii, mitkä prosessilouhintatekniikat soveltuvat suurten potilasdatajoukkojen mallintamiseen ja kuinka saatuja tuloksia voidaan analysoida terveydenhuollon johtamisen näkökulmasta.

Prosessilouhinta osoittautui arvokkaaksi työkaluksi suurten potilasdatajoukkojen mallintamiseen. Tulokset osoittavat, että lasten ja nuorten päivystyspotilaiden hoitopolut ovat usein samankaltaisia, ja hoitoon tarvitaan henkilöresurssien lisäksi vain vähän diagnostisia ja hoidollisia resursseja. Poikkeuksena ovat akuutit ESI 1 -luokan potilaat, jotka tarvitsevat huomattavia resursseja ja seuraavat yksilöllisiä hoitopolkuja. Kokonaiskuvassa ESI 2, 3, 4 ja 5 -luokkien välillä havaittu samankaltaisuus viittaa siihen, että ESI triage -järjestelmä ei toimi tarkoituksenmukaisella tavalla.

Avainsanat Prosessilouhinta, Lasten päivystys, Terveydenhuollon johtaminen

Preface

I would like to thank my advisor, Docent Miika Koskinen, who gave me the opportunity to undertake this project and has supported me throughout the research process. I would also like to thank my supervisor, Professor Lauri Saarinen, who placed his trust in the project and gave me the freedom to shape it according to my vision. Finally, I would like to thank everyone at the HUS organization who has been involved in supporting and assisting with this thesis project.

What an unforgettable journey it has been. About five years ago, I ended up at Aalto University almost by accident. It was a flash of an idea that turned into a long-lasting commitment. Now, the journey is coming to an end. While my studies will continue for a few more years in the field of medicine, my time at Aalto has been experienced. With its ups and downs, I have always made sure to shape it into something that reflects my own vision.

This thesis is my answer for evergoing guestions for the path I have chosen. One might be too much of an engineer to be a medical doctor and too much of a medical doctor to be an engineer. Yet, by combining these fields of science, there might be something unique in practice.

One might have a goal, but without a process, it is just a dream.

Helsinki, 13 February 2025

Tuomo Antikainen

Contents

Abbreviations

BPR	Business Process Redesign
EBM	Evidence-Based Management
HUS	HUS Helsinki University Hospital
BI	Business Intelligence
VBHC	Value-Based Healthcare
ED	Emergency Department
ICT	Information and Communication Technologies
XES	eXtensible Event Stream
XML	eXtensible Markup Language
BPMN	Business Process Model and Notation
EPC	Event-driven Process Chain Flowchart
BPM	Business Process Management
YAWL	Yet Another Workflow Language
DFG	The Directly-Follows Graph
BPM	Business Process Management
iDHM	interactive Data-aware Heuristics Miner
DFM	Directly Follows Model Miner
PM4PY	A process mining library for Python
IDE	Integrated Development Environments
CSV	Comma-Separated Values
NumPu	Numerical Python
Pandas	Python Data Analysis Library
NCH	New Children Hospital
ESI	Emergency Severity Index
EMS	Emergency Medical Services
LOS	Length Of Stay
GP	Genetic Programming
HIS	Hospital information system
GQFI	Goal-Question-Feature-Indicator
ECG	Electrocardiography
POC	Point-Of-Care test

1 Introduction

Personalized patient care processes that aim to diagnose, treat, and prevent diseases to improve the general health of patients form the core of healthcare [1]. These actions are carried out by resources, including healthcare professionals, diagnostic and treatment devices, infrastructure, and supporting systems. Together, these elements make up the complex network known as the healthcare system.

One of the biggest challenges facing healthcare systems globally in the 21st century has been overcrowding, where demand exceeds supply [2]. This issue has forced a rethinking of how to create more effective and efficient healthcare delivery. Methods like Lean, Business Process Redesign (BPR), and Evidence-Based Management (EBM) have been utilized in efforts to streamline frontline processes of healthcare [3, 4, 5]. However, it is widely recognized that this task is challenging due to the dynamic, complex, and ad-hoc nature of healthcare processes [6]. Recently, a relatively new approach, Process Mining, has emerged as a potential solution to address these challenges.

1.1 Background and Motivation

During the past decades, healthcare services at HUS Helsinki University Hospital (HUS) have generated substantial amounts of data, which hold significant economic and practical value for EBM. Effectively leveraging this data through advanced analytical methods has the potential to reduce cost pressures and enhance operational efficiency. However, traditional business intelligence (BI) techniques are often inadequate for fully exploiting data to improve long-term healthcare processes. Key aspects such as cost efficiency, resource allocation, and medical effectiveness, which are central to Value-Based Healthcare (VBHC), remain challenging to optimize using conventional approaches [7].

The heterogeneity of the population within a single disease or across different stages of a disease introduces substantial variation in care processes. Notably, the variability of processes in healthcare far exceeds that of manufacturing or other service industries [6]. This highlights the critical need for robust process analysis methodologies to visualize, analyze, and manage the intricate dynamics of healthcare processes, such as patient pathways. To enhance care efficiency and control costs, it is necessary to understand how variability driven by diverse patient characteristics and needs affects these pathways. Equally essential is identifying and understanding how inefficiencies within service processes, such as resource shortages and high workloads, affect pathways and contribute to failure demand. Additionally, natural factors, such as time of day, day of the week, and holiday seasons, must be acknowledge, as they contribute to seasonality, cycles, and trends.

Process mining offers a robust, data-driven approach for addressing these distinguishing characteristics and challenges in healthcare. While traditional BI techniques may struggle to fully capture the complexity and dynamics of healthcare operations, process mining enables a more detailed visualization and analysis of these processes, uncovering insights that might otherwise remain hidden. In doing so, it shifts decision making from relying on planned processes to being guided by actual process data, thereby enhancing opportunities for evidence-based operations management.

1.2 Research Questions and Objectives

This master's thesis project focuses on exploring how process mining can be applied to map patient pathways in the pediatric and adolescent emergency departments (ED) at HUS. By doing so, the aim is to gain a better understanding of how these emergency departments operate. The study has three main objectives. The first objective is to establish the capability and expertise for HUS pediatrics to leverage process mining. As process mining has never been used at HUS pediatrics, a platform for its application needs to be developed from scratch. The goal is not only to create an environment for this project but also to ensure that the platform and acquired knowledge expand for widespread use at HUS in the future.

The second objective is to investigate how current process mining tools can be utilized with large healthcare datasets in the ED setting. The study uses three full years of data from 2021 to 2023 from pediatric and adolescent ED. To date, no studies have been conducted on pediatric or general emergency departments at this scale of patient volume, making it essential to investigate whether meaningful insights can be derived from such large datasets. In parallel, this study provides foundational groundwork and serves as a pioneering example for future research in this area.

The third objective is to produce process models from which conclusions can be drawn about the operations of the pediatric and adolescent ED. The study focuses specifically on resource allocation, pathway variability, and failure demand in relation to mistriage. Triage is a system used in the ED to prioritize patients according to the urgency of their condition, ensuring that they receive timely and appropriate care [8]. By examining resource allocation, pathway variability, and failure demand, the objective is to better understand ED operations and identify possible areas that are not functioning as intended or could be made more efficient.

Based on these objectives, this thesis aims to address three research questions:

- To what extent are existing process mining tools capable of effectively handling large scale healthcare datasets?
- Can these tools be utilized to create meaningful patient pathway process models from pediatric and adolescent emergency department data?
- What patterns of resource allocation and pathway variability can be observed in pediatric and adolescent ED, and is there a failure demand?

1.3 Structure of the Thesis

The structure of this thesis is as follows. Chapter 2 introduces the foundations of process mining, providing the theoretical basis for the study. Chapter 3 explores the application of process mining within the healthcare context. Chapter 4 offers an overview of HUS, focusing on its pediatric and adolescent emergency departments and the ESI triage system employed. Chapter 5 reviews key literature on the use of process mining in pediatric and adolescent emergency departments, as well as in broader emergency department settings. Chapter 6 describes the data utilized in this study and outlines the methodology, including data preprocessing and process discovery techniques. Chapter 7 presents the findings, while Chapter 8 interprets these results, discusses the study's limitations, and evaluates the applicability of process mining in this context. Finally, Chapter 9 summarizes the study and its contributions.

2 Foundations of Process Mining

Many industries have experienced a significant increase in data availability in the 21st century. This is a consequence of major shift towards digitalization in societies and industries [9]. Rapidly evolving information and communication technologies (ICT), such as online tools, social networks, and smart sensors generate vast amount of heterogeneous data every day [10]. The term "Big Data" is frequently used to describe this situation [11]. In the last decade, industries have become more aware of the opportunities of data collection [12]. This awareness has prompted many businesses to implement comprehensive data collection strategies and establish data warehouses [13]. However, the primary objective for most businesses is not the accumulation of data, but the transformation of data into actual value using proper analysis. This concept is known as Business Intelligence (BI) [14]. While classical data mining has been a core component of BI, the growing curiosity and demand for innovative data analysis methods have integrated additional techniques into BI practice. One of these methods is process mining.

Process mining is a process management technique which is used to discover, monitor, and enhance real processes. It combines a family of retrospective analysis methods by utilizing event log data extracted from ICTs. By utilizing this data-centric approach, process mining serves as a connecting link between conventional model-based process analysis methods, such as process simulation, and data-driven approaches like data mining and machine learning [15]. Additionally, it can also be considered a overlap of data science and process science. Originating as a new sub-field of data mining in the early 21st century, process mining has undergo major evolution to become more routinely used separate analysis method. As a result, it has reduced the reliance on intuition and increased data-driven decision making in operations management.

Process models, which were introduced for process simulation over 50 years ago, remain at the heart of process mining. These models illustrates the dependencies between actions within a process. While process models have been extensively used across various industries to enhance efficiency, improve quality, and ensure compliance, classical process simulation struggles with a significant challenge: the complexity of real-world processes. Manually constructed process models often present an overly simplistic and optimistic view of actual processes [16]. A common error is to simulate what is intended to occur rather than reflecting what is happening. This can lead to decisions based on false assumptions. In contrast, process mining is designed to eliminate assumption-based errors and provide an accurate, data-driven representation of actual processes. This aims to offer a more objective view of processes and to unlock new possibilities for utilizing process models.

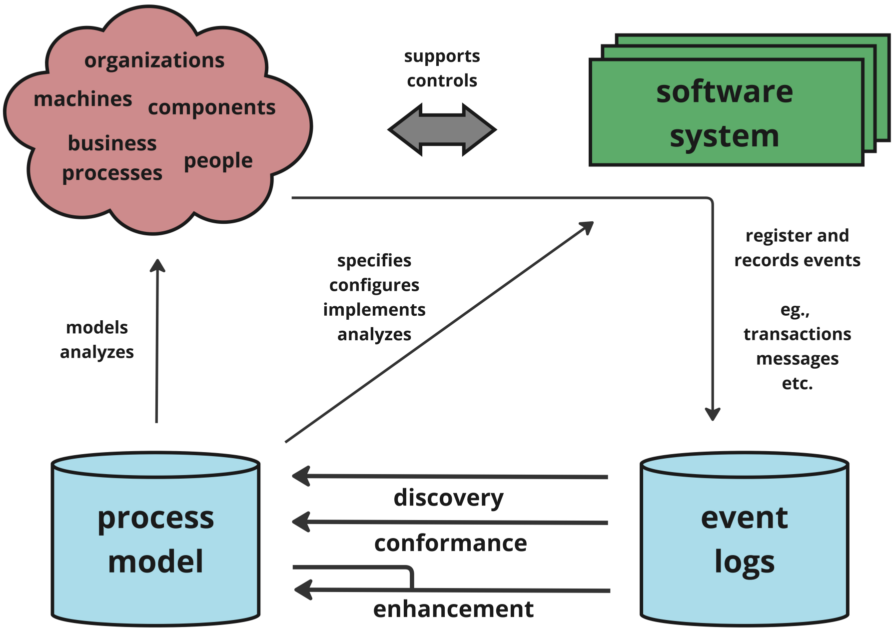

Figure 1: Discovery, Conformance, and Enhancement. The three basic types of process mining (based on [15]).

Figure 1 illustrates the three main types of process mining techniques, discovery, conformance, and enhancement, positioned within the process mining cycle. This cycle comprises three primary phases: data extraction (event log registration), process mining, and analysis. This chapter offers a holistic overview of these phases and the fundamental principles of process mining. It begins with a detailed examination of event logs, followed by an exploration of process mining techniques and perspectives. Subsequently, the practical value and challenges associated with process mining are addressed. The chapter concludes with a discussion on miners, process mining tools, and the distinctions between process mining and data mining. The objective is to explain the fundamental structure of process mining and demonstrate how it can be leveraged to extract meaningful insights and real value from Big Data.

2.1 Event Log Data

Process mining utilizes event log data for analysis. Modern ICT systems register data in many forms including event logs. Event logs have a structure tailored to enable the representation of key aspects of events occurring across various iterations of a specific process [17]. In simplified terms, an event log is a file created by an information system that contains data on the execution of a process. In principle, event logs are formulated in the way each case is described by one or more events and each event has one or more attributes. More precisely, following relationships must hold [15]:

- Each event in the event log relates precisely to single case.
- Each case in the event log relates precisely to single process.
- Additional information may exists whereupon it is attached to cases or events as attribute.

To get better intuition about event logs, simplified hospital ICT system can be used as an example. During hospitalization of a patient, every task performed on them is registered to ICT either automatically or manually. This include for instance doctor's notes, procedures, laboratory tests and prescribed medications. Each registration creates an individual event in an event log. This leads to the event log being filled with events that characterize patient cases, also called patient pathways or traces. In process mining, these cases combined can be viewed and researched as a single patient flow process.

Table 1 illustrates an example of event log described above. It shows the main information categories that need to be recorded in an event log used for process mining. Three main categories are included which are case id, event id, and attributes [15]. Case id and event id are on theory level self-explanatory. Case id is a unique identifier for all events that relate to the same case. Event id on the other hand is a unique identifier for every event. However, in practice, event log may include many case categories which divide events differently. In this situation, it is important to choose case id in that way it is in line with research question.

To enable proper process mining analysis, besides case id and event id, event logs must include timestamp and activity attribute categories [18]. Timestamps allow correct sequencing of events and are also vital for measuring process performance. Activities illustrate steps or operations that make up a larger process and are fundamental part of analysis. Above mentioned event log categories are enough to perform process mining analysis. However, to get more insight, additional event attribute categories are usually introduced to analysis. These can be for instance used resources (e.g. people, machines), locations, and costs.

In addition to the data itself, data editing is an important factor for successful process mining. Real life processes rarely generate data that is suitable for process mining as it is[17]. Hence heterogeneous data from many sources must be standardized and cleaned from incorrect logs. The de facto standard event log format for process mining is eXtensible Event Stream (XES) [15, 19]. XES design allows custom attribute setting to support different needs of users. As eXtensible Markup Language (XML) based format, there are along with process mining tools also various sets of other software tools to work with.

Table 1: Example of event log. The table does not contain real patient data but demonstrates simplified hospital emergency department event log data.

Event id	Case id	Attributes			
		Activity	**Timestamp**	**Department**	**...**
491083	1	Triage	2023-02-03T13:25:36	ED	...
491084	1	Blood profile	2023-02-03T13:32:21	Laboratory	...
491085	1	Lung x-ray	2023-02-03T13:47:02	Radiology	...
491086	1	Medication	2023-02-03T14:02:19	ED	...
922037	2	Triage	2023-02-03T13:33:16	ED	...
922038	2	Suturation	2023-02-03T14:04:59	ED	...
259634	3	Triage	2023-02-03T14:11:45	ED	...
259635	3	Blood profile	2023-02-03T14:18:04	Laboratory	...
259636	3	Medication	2023-02-03T14:56:21	ED	...
305952	4	Triage	2023-02-03T14:40:19	ED	...
305953	4	Lower leg x-ray	2023-02-03T15:02:02	Radiology	...
305954	4	Plaster casting	2023-02-03T15:48:56	ED	...
329513	5	Triage	2023-02-03T15:03:41	ED	...
...	...	...	...	...	...

2.2 Three types

Process mining is typically separated into three distinguished types which are discovery, conformance, and enhancement. Although all of these are used to analyze, monitor, and improve processes, each type has its own focus and objectives.

Process discovery is regarded as one of the most challenging aspects in the field of process mining [15]. It involves extracting a process model from event log data without advance knowledge of the process. [20]. Typically approached from a control-flow perspective, process discovery aims to uncover and visualize the sequences and dependencies of activities within a process. A notable advantage of this method is its capacity to develop data-driven process models that are both objective and accurate, while remaining simple and easy for users to understand. This capability to discover processes that were previously difficult to model has raised great interest in many business sectors [21]. Consequently, it has become the most widely utilized type of process mining. Figure 2 presents an example of BPMN process flow chart which illustrates these sequences and dependencies of activities in hospital ED. However, in real life processes are rarely this simple which creates difficulties and limitations for process discovery. In addition for BPMN, Petri nets, EPC, Causal nets, DFG, and Process trees are commonly used for visualization in process discovery [15, 22, 23, 24].

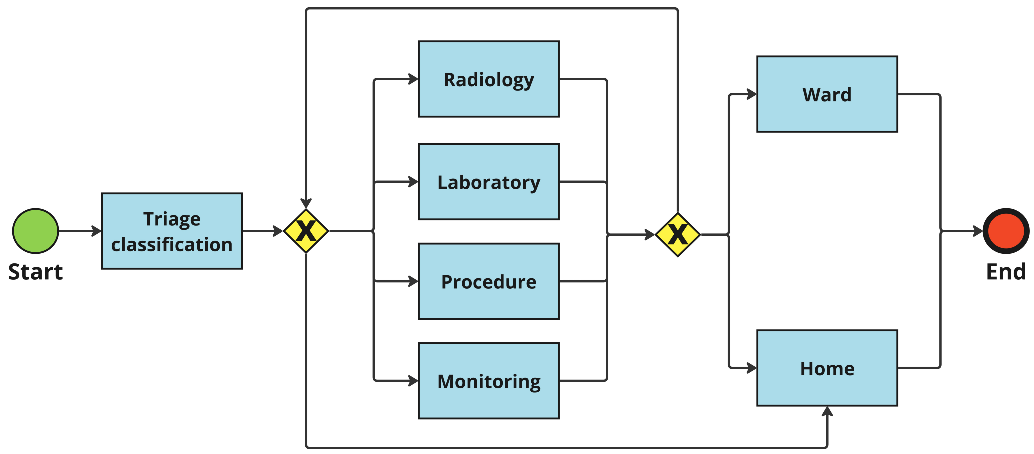

Figure 2: Simplified hospital emergency department BPMN process model.

Conformance checking is used to detect inconsistencies between a process models and reality [25]. An opted model, either hand constructed or discovered, is assessed on whether it conforms to registered behavior in event logs. If inconsistencies are observed, these are identified and analyzed in more depth. This type of analysis is increasingly important in modern process management [26]. Frequently changing dynamic processes, such as those found in healthcare, are challenging to model with traditional Business Process Management (BPM) tools. Reliable models being important for process management, conformance checking can be applied to evaluate the performance of these models.

Process enhancement is used to improve and extend existing process models utilizing information from event logs [27]. In a situation where an inaccurate process model is found, it can be adjusted using the diagnostic information derived from the alignment of the model and the log. Moreover, an accurate process model can be reinforced with additional attribute information received from event log. This information can be for instance frequencies, resources, costs, bottlenecks, and execution times [15].

2.3 Perspectives

Event logs typically contain an extensive repository of information in attributes. This data is feasible for expanding the modeling perspective beyond the control flow perspective. Commonly used additional perspectives are organizational perspective, time and probability perspective, and case perspective. The organizational perspective is utilized to study process resource allocation, time and probability perspectives to analyze process performance, and the case perspective to examine process dependency rules [15]. While the process flow perspective remains the most widely used, alternative mining perspectives provide valuable insights that deepen the understanding of processes. These perspectives can be directly applied in process discovery or leveraged to enhance existing process models.

Organizational perspective focuses on utilizing resource attributes present in most event logs. Typically, this requires human resources but other means such as machinery may be included as well. Organizational perspective is divided further into three sub perspectives which are social network, organizational structure, and resource behavior perspectives [15]. Social network perspective concentrates on mine relationships between resources. It is employed to create visualization of social networks of resources. Relationships are weighed to reflect their importance in the process. The concept of importance can be defined in multiple ways to suit specific analysis. The organizational structure perspective mines relationships between resources and activities, showing how regularly certain activities are carried out by specific resources. It helps identify similar resources within a process. Resource behavior perspective analyzes resource performance by examining the frequencies and duration of activities performed by individual resources. The outcome provides insight into the workload of individual resources.

Time and probability perspective introduces timing and frequency considerations into process mining. It utilizes timestamps associated with activities in event log. Time and probability perspective is an essential component of process performance mining. It allows examination of performance from multiple sub perspectives including waiting and service time, flow time, bottleneck, and frequencies and utilization perspectives [15]. The perspective of waiting and service time focuses on analyzing the transition times between sequential activities and the duration of each activity. This perspective is valuable on its own but also for bottleneck analysis, as it helps identify areas where delays occur and where resources may be underutilized. The bottleneck perspective specifically aims to detect and analyze bottlenecks within a process. Bottlenecks are activities within a process that disrupt the flow and lead to underutilization of downstream resources [28]. Flow time perspective analyzes the time taken for cases to move through a process. It can be used to calculate the flow time for specific end-to-end process routes or between any two arbitrary points within the process. This perspective is valuable for comparing different process routes and identifying variations in flow times. Additionally, it can be used to determine how often specific parts of the process meet their target times. The frequencies and utilization perspective is used to calculate routing probabilities, which represent the conditional probability of the next event given previous observations. This approach assesses the likelihood that, if a case reaches activity X in the process, it will subsequently proceed to activity Y. For instance, in the process flow chart shown in Figure 2, this could refer to the probability that case will move on to radiology imaging after the triage classification. Probabilities are derived from route frequencies observed in cases and can be calculated between two contiguous activities or across extended sequences of activities.

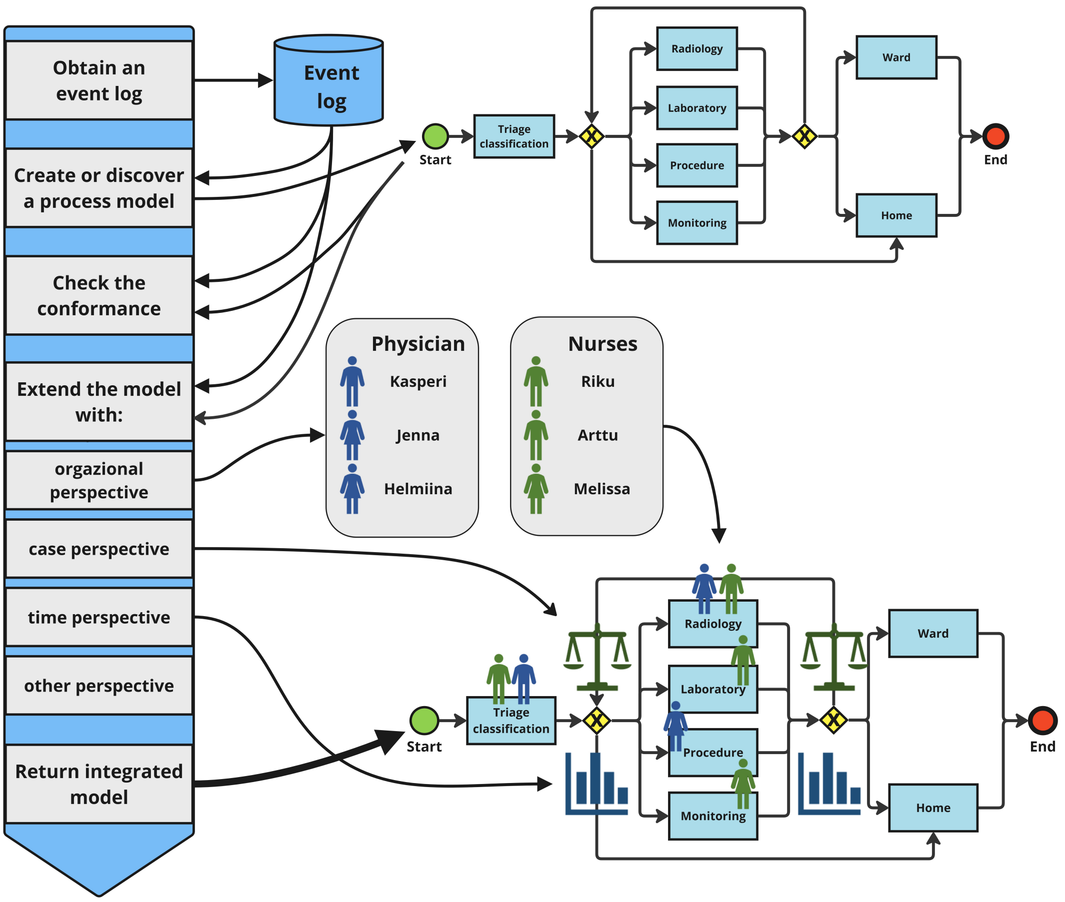

Figure 3: Fully integrated process model including organizational, time, and case perspectives (based partly on [15]).

The case perspective, also referred to as decision mining, aims to identify rules that explain the paths taken by individual cases [29, 30]. The core principle is to identify process decision points and the factors that affect decision-making. Process decision points are junctures in a process where multiple possible directions can be taken. For instance, in the BPMN process flow chart shown in Figure 2, decision points are represented by yellow diamond shapes. At these points, decisions – whether trivial or complex – are made to direct individual cases along specific paths. For instance, a trivial decision might involve a binary factor, while a complex decision could involve multiple factors. The case perspective utilizes case attributes to identify these factors and uncover insights into decision-making throughout the process.

The best understanding of the analyzed process is achieved by combining multiple process mining perspectives. Figure 3 outlines an example of a process model where organizational perspective, time and probability perspective, and case perspectives have been merged with process with standard control flow perspective. Although combining all these perspectives allows to perform an extensive analysis of the process, it is important to note that the list of perspectives presented in this section is not complete. As process mining continues to evolve, new methods and perspectives for its application are constantly being developed.

2.4 Practical Value

Process mining has evolved considerably in the last two decades. It was initially developed as an academic research method but was quickly adapted into operations management. Within the context of BI, the significance of process mining lies in its capacity to transform academically proven methods into practical value. Process mining is already applied in various forms and from different perspectives within BI. Its implementation can vary widely, with the desired outcomes differing significantly across studies and analyses. Nevertheless, the practical value delivered by process mining can be broadly classified into six key areas [31]. These areas are listed in Table 2.

Insight provision forms the basis of process mining. It relies heavily on process discovery techniques. This aspect of process mining uncovers hidden patterns and details within processes, providing a clear and comprehensive understanding of how processes are performed. Frequently, the variability within processes is much greater than initially expected or planned [32]. Identifying this variability between expected or planned processes and actual processes holds significant value for organizations, as it highlights sources of waste and mismanagement.

Variability identification is essential for a holistic understanding of processes. As mentioned, variability within processes can be substantial. While insight provision delivers information about variability between expected or planned processes and actual processes, variability identification extends this analysis to variability between processes over time. Process discovery, process performance, and process conformance techniques can be utilized to analyze how different instances of a process vary in terms of duration, sequence, and outcome. By identifying these variations and causes, organisations can implement measures to standardize processes.

Performance improvement is attained through the analysis of process performance attributes. This involves employing various performance perspectives to identify areas where enhancements can be made. For example, by examining bottlenecks, flow times, and resource utilization, organisations can implement targeted interventions to streamline operations, reduce cycle times, and optimize overall performance. Also, other performance factors such as cost, or customer satisfaction can be utilized.

Conformance assurance is achieved with conformance checking. This helps organisations to ensure that processes consistently operate within the limits defined by managers and regulatory authorities over time. Any deviations over the defined limits detected through process mining could reveal issues such as fraud, malpractice, risks, or inefficiencies [31]. Moreover, addressing these deviations can significantly enhance overall process compliance and operational effectiveness.

PV1	Insight Provision
PV2	Performance Improvement
PV3	Conformance Assurance
PV4	Variability Detection
PV5	Reliability Enhancement
PV6	Prediction

Table 2: Practical values of process mining.

Reliability enhancement is achieved by detecting the root causes of encountered failures within systems or processes. To accomplish this, process discovery, process performance, and conformance checking techniques are employed, utilizing historical data related to system or process failures. This fault diagnostic enables organisations to address the underlying causes of failures, ultimately improving process reliability and leading to more dependable and resilient operations.

Prediction holds perhaps the most significant potential for the future of process mining. The primary goal is to predict how specific cases within the process will behave in the future. This enables organizations to make predictive process adjustments enhancing the ability to meet both qualitative and quantitative requirements for upcoming cases [33]. Currently, with process discovery techniques and with decision mining perspective, it is possible to retrospectively analyze processes and the factors that have influenced decision making within them. These methods also enable statistical predictions, forecasting how specific cases with certain attributes are likely to progress through the process.

2.5 Challenges

Although process mining has evolved to be a widely used analysis method, there remain fundamental challenges that must be addressed. The Process Mining Manifesto identifies eleven critical challenges [21], which are listed in Table 3. Of these, challenges C1, C2, C4, C5, C6, C10, and C11 are discussed in more detail below.

Challenge 1, finding, merging, and cleaning event data, and Challenge 2, dealing with complex event logs with diverse characteristics, are both related to data extraction and preprocessing. These tasks are typically the initial and most critical steps in process mining. This is especially significant when dealing with big data. Although the increasing trend of data collection has facilitated the extensive application of process mining it also means that collected data is from a wide variety of sources. This causes problems when collected data has varying identifiers, different levels of granularity, and insufficiencies. For instance, one system might use a name and system generated code to identify an individual, while another system might rely on the social security number of people. Similarly, one system might record timestamps with precision to the second, while another system may log them hourly or not record them at all. Overall, the attributes recorded by different systems can vary significantly. Consequently, merging data can sometimes be highly challenging.

C1	Finding, Merging, and Cleaning Event Data
C2	Dealing with Complex Event Logs with Diverse Characteristics
C3	Creating Representative Benchmarks
C4	Dealing with Concept Drift
C5	Improving the Representational Bias Used for Process Discovery
C6	Balancing between Quality Criteria
C7	Cross-Organizational Mining
C8	Providing Operational Support
C9	Combining Process Mining with other Types of Analysis
C10	Improving Usability for Non-experts
C11	Improving Understandability for Non-experts

Table 3: Key challenges in process mining.

After data merging, there are also significant variability in event log characteristics, which can affect suitability for process mining. Very large event logs present challenges in mining, while smaller event logs frequently lack sufficient data, making it difficult to derive reliable conclusions. However, the absolute number of cases and events alone does not determine the difficulty of analyzing an event log. Factors such as the mean number of events per case, the analogy between cases, the count of unique events, and the diversity of paths significantly impact the complexity of the analysis. As outlined in the Process Mining Manifesto, an event log with 1000 cases, each averaging 10 events and exhibiting minimal path variation, is generally easier to analyze than an event log with 100 cases, each averaging 100 events and following unique paths, despite both logs having approximately the same total size [21]. Currently, a significant challenge is the lack of a straightforward and reliable method for identifying which event logs are suitable for process mining, aside from using time consuming a trial-and-error approach.

Challenge 4, dealing with concept drift, involves addressing the dynamic evolution of a process. For instance, this can manifest through periodic or seasonal variations within the process. In such cases, a single process model is insufficient to accurately represent the true process over time. Concept drift can be identified by dividing the event log into smaller sections and analyzing the distinct patterns within each segment. However, this requires a substantial amount of event log data collected over an extended period. This is because event log collected over a short period may not yet reflect periodic or seasonal changes clearly enough to distinguish them from normal process variability. In simple terms, single analysis conducted over a long period may overlook concept drift, while an analysis conducted over a short period may fail to detect it. Although variability detection analyses are available, it currently needs to be conducted separately. This means that the analyst must have a clear understanding of what concept drift is. In the future, techniques must be developed that account for this automatically and more efficiently, as identifying and analyzing concept drift is essential in process management.

Challenge 5, improving the representational bias used for process discovery, and Challenge 6, balancing between quality criteria, addressing the discrepancy between mined process models and reality. In an ideal situation, the process model would perfectly represent the process under analysis, but this is unfeasible. There are two significant reasons for this: the limitations of process modeling languages and the incompleteness and noise in event logs. Visualizations for process models are created using modeling languages such as DFG, Petri nets, and BPMN. It is necessary to recognize that these graphical representations differ from the computational models generated through various process mining techniques. Each of these languages encompasses several implicit assumptions and may not be suitable for visualizing all types of processes. Therefore, in process mining, analysis should not be restricted by a preferred model language but rather by selecting the best language based on the specific needs of the process.

The incompleteness of event logs imply that they rarely represent the actual process in full, but rather only a part of it. Process models, by design, allow for exponential growth in the number of possible paths as the number of activities increases. However, the likelihood of each path varies, with some having a extremely low probability of occurring. As a result, the probability that the event log contains all possible paths in the process is practically zero. This can be clarified with an example given in the Process Mining Manifesto [21]. Let us consider a process consisting of 10 activities that can be executed in parallel, alongside a corresponding event log containing data on 10 000 cases. In this model with 10 concurrent activities, the total number of possible interleavings is 10! = 3 628 800. Therefore, it is not feasible for every interleaving to be represented in the event log, as the number of cases is far fewer than the possible paths. Even if the case count were increased to millions, capturing all paths would remain highly unlikely. This introduces challenge known as noise.

Noise refers to less frequent paths that deviate from the main flow of the process. It can be caused by low frequency correct process behavior, errors in the process, or data transcription mistakes. The challenge is that distinguishing between these causes is highly complex [34]. Furthermore, due to the incompleteness of event logs, it is impossible to definitively ascertain whether these low frequency paths are indeed more probable than certain paths that are missing from the event log. Consequently, creating a model that accurately reflects the investigated process is generally infeasible. To assess process models, four quality criteria have been established: fitness, simplicity, precision, and generalization [15].

Fitness quantifies how accurately the model reflects the event log used. A model with perfect fitness contains every path present in the event log from start to finish. Simplicity characterizes how well the model can explain the event log while preserving a simple design. Although these two criteria are often viewed as opposites, relying solely on them is inadequate for comprehensive model evaluation. This is because it is possible to create a model that achieves perfect fitness while also being extremely simple. However, such a model also represents many other potential event logs associated with the given activities. Consequently, this leads to a lack of specificity regarding the analyzed event log, a phenomenon commonly referred to as overfitting. Precision evaluates how specific the model is to the analyzed event log. Yet, because event logs are incomplete, striving for perfect precision often results in underfitting, in which the model fails to adequately represent the actual process. To avoid this, generalization, which can be considered as the opposite of precision, assesses how well the model accounts for possible process behaviors that are missing from the event log. By balancing these four quality criteria, the impact of event log incompleteness and noise on process mining can be reduced. Nevertheless, achieving this balance is extremely challenging. Future developments should focus on creating more efficient techniques to balance these four criteria.

Challenge 10, improving usability for non-experts, and Challenge 11, improving understandability for non-experts, are related to the usability of process mining. For process mining to integrate seamlessly into everyday operational support, its usability must be enhanced to be more an intuitive and user friendly. Currently, leveraging process mining effectively demands a high level of expertise in data science. While some commercial process mining software provides more user-friendly interfaces such as Disco and Celonis, this often comes at the expense of advanced algorithmic capabilities. A future challenge is to integrate sophisticated algorithms within intuitive and half automated interface, ensuring effective process mining with a lower entry barrier. Similarly, attention must be given to process mining results. Even if the outcomes appear clear, there is a risk of drawing incorrect conclusions from them. A significant risk related to this is making overly broad conclusions from event logs collected over a short period time or from overly simplified models derived from complex event logs. Current process mining algorithms do not generally alert users of models derived from such event logs, that may suffer from low fitness or overfitting. Thus, attention must be given to clear methods for presenting process mining results and ensuring trustworthiness.

2.6 Miners

Process mining relies on algorithms specifically designed to extract insights from event log data. These algorithms are grounded in mathematical principles, including graph theory, statistics, and optimization, and are translated into computational processes through various programming languages [15]. A diverse array of process mining algorithms has been developed to support various types and perspectives of analysis. Selecting the appropriate algorithm is essential for achieving the desired outcomes.

The first developed classical algorithms, that are commonly referred to as miners, support process discovery. By analyzing the sequences of events in the log, the aim is to generate accurate but readable visualization of the process. For process discovery there are for example Directly Follow Model Miner, Alpha Miner, Heuristic Miner and Fuzzy Miner [35, 36, 37, 38]. Second group of algorithms are developed support conformance checking. These aim to check and detect inconsistencies between a process model and reality. For conformance checking there are for example Footprints and Alignments algorithms [15, 39]. Process enhancement includes many perspectives of process mining. Some utilize specific algorithms, while others are analyzed through a combination of existing algorithms and extensions. For example organizational perspective includes Social Network algorithm and Organizational Structure algorithm [40, 41]. On the other hand, the time perspective can be implemented using existing algorithms like the Inductive Miner and Fuzzy Miner. A good example of an algorithm that integrates multiple existing techniques is the interactive Data-aware Heuristics Miner (iDHM). iDHM combines elements from the Alpha, Heuristic, and Fuzzy Miners, and is capable of handling process discovery, conformance checking, as well as the time perspective [42, 43].

Process discovery miners are crucial for the purposes of this research. Alpha Miner is recognized as the first foundational process mining algorithm. It utilizes Petri nets as a visualization tool and employs a straightforward pattern recognition method to generate process models. Essentially, the Alpha Miner assumes a causal dependency from two arbitrary actions, X and Y, if the event log contains cases where X is followed by Y, with no instances of Y preceding X [36]. In this scenario, a direct connection is established from event X to event Y. While the algorithm is theoretically simple and efficient, in practice it is unable to handle noise, concurrency, and the incompleteness of event logs. As a result, it is rarely used with real life event logs.

Heuristic Miner was developed to address the limitations of Alpha Miner. A significant improvement in Heuristic Miner is its ability to account for path frequencies. Noisy event logs often include instances where two arbitrary actions, X and Y, are recorded such that X follows Y, as well as instances where Y follows X. While Alpha Miner is unable to effectively process these situations, Heuristic Miner utilizes a dependency graph to address this issue [37]. The dependency relation $X \implies Y$ between two arbitrary actions X and Y in an event log is calculated

$$X \implies Y = \frac{|X > Y| - |Y > X|}{|X > Y| + |Y > X| + 1} \tag{1}$$

where $|X > Y|$ is count of instances where X precedes Y and $|Y > X|$ is count of instances where Y precedes X. The dependency relation is therefore always between -1 and 1. Heuristic Miner calculates the dependency relation between all actions present in the event log. This approach helps reduce the impact of noise in the event log by establishing threshold for the dependency between actions that are included in the process model. Heuristic Miner is also capable of calculating long distance dependencies, enabling better concurrency and loop modeling [44, 37]. For visualization, Heuristic Miner provides several options, including dependency graphs, Petri Nets, and Causal Nets. A drawback of the Heuristic Miner is that it requires careful parameter setting to ensure optimal results. Finding the right parameters can be highly challenging, and even with fine adjusting, there is a risk of overfitting or producing overly complex models.

The Fuzzy Miner was developed to address the challenges of complex event logs. While Heuristic Miner can manage noise, it tends to overfit or create excessively complex process models from event logs with a substantial number of path variations. To address these shortcomings, Fuzzy Miner leverages four core principles [38]:

- **Aggregation:** clustering less significant information to reduce complexity of model.
- **Abstraction:** omitting insignificant information to reduce complexity of model.
- **Emphasis:** highlighting significant information to clarify model.
- **Customization:** allowing adjustments to the level of detail to meet the specific needs of the model.

Transforming these principles into an actual algorithm requires the definition of two fundamental metrics: significance and correlation. Significance measures the relative importance of behavior within a process [38]. It can be assessed for both individual events and the paths that connect them. Relativity in this case suggests that the criteria for significance can be flexibly defined and adapted according to the specific goals of the process model. As a straightforward approach, frequency is a commonly used measure of significance. In this context, events and paths recorded in the event log that occur more frequently are considered more significant than those recorded less often. Correlation measures the degree of relationship between two events that occur sequentially [38]. Unlike significance, it can only be measured for events. There are various methods to measure correlation. However, the fundamental principle is that for a correlation to exist, events must share a significant number of process instances from the event log.

With these two metrics, the functionality of the Fuzzy Miner algorithm can be simplified as follows [38]:

- High-significance behavior is preserved in the model.
- Less significant yet highly correlated behavior is clustered within the model.
- Less significant and weakly correlated behavior is removed from the model.

Table 4: Summary of the strengths and weaknesses of four process mining algorithms.

Miner	Strengths	Weaknesses
Alpha Miner	Simple, easy to understand, foundational algorithm	Requires a complete log, Struggles with: noise, concurrency, loops
Heuristic Miner	Handles noise, uses frequency, detects concurrency	Requires careful tuning, can overfit, may become complex with large logs
Fuzzy Miner	Handles noisy and complex processes, visually simplifies	Less precise, less formal, hard to interpret finer details
DFM Miner	Simplicity, vagueness, scalability	Oversimplifies, struggles with: concurrency, noise, incompleteness of event log

This enables comprehensible representation of complex processes. Additionally, it allows for extensive customization through the selection of significance measures, as well as the adjustment of thresholds for both significance and correlation. A drawback of Fuzzy Miner is that the visualization method it employs, the Fuzzy net, is not as formal as, for example, a Petri net. Additionally, the algorithm is not well suited for precise modeling, as its one primary objective is to simplify complex processes.

Directly Follows Model Miner (DFM) is the simplest process discovery algorithms in process mining. It builds a process model by mapping out activities in the exact order they follow each other in the event log, without considering more complex relationships like concurrency or loops [35]. DFM visualizes processes using Directly Follows Graphs (DFG), which are simpler than higher-level visualization languages, showing only the flow direction and frequency between activities.

The simplicity of the DFM algorithm provides distinct strengths as well as weaknesses. Its key advantages include model clarity, vagueness, and scalability, while it may suffer from oversimplification [45, 46]. Model clarity refers to its capacity to generate easily interpretable models from complex event logs, even when higher-level visualization languages are considered un-interpretable. Vagueness means that DFM can express relationships that other miners cannot interpret with precise clarity. Furthermore, DFM can efficiently process event logs containing millions of events, a scale at which other miners often struggle to perform.

On the downside, DFM struggles to manage noise, incompleteness, and other exceptions within the event log. It either includes all transitions from the event log or relies on frequency-based filtering, both of which present challenges when applied to real-world processes [46]. Additionally, since DFM does not account for concurrency, the generated process model can only reflect immediate transitions and poorly handles long distance dependencies. Table 4 provides a concise summary of the strengths and weaknesses of the miners discussed in the preceding paragraph.

2.7 Tools

In the last two decades, many process mining software have become available to conduct process mining analysis. These include open source free to use software such as ProM and Apromore and commercial software such as DISCO and CELONIS. All these share core principles of process mining. However, the purpose of use and extent of the analysis methods differ. In addition to the aforementioned, a process mining library for Python (PM4PY) was introduced in year 2019 [47]. It is a novel programming language base process mining tool which has gain much popularity within a few years.

ProM has now been the de facto process mining tool for 20 years. It revolutionized process mining field after launching 2004 [15]. Before that, all few available process mining software were supporting single technique and perspective of process mining. The main idea of the ProM was to create software that brought together various types of process mining techniques, perspectives, and algorithms [48]. As a result, ProM expanded vastly the field of possibilities in process mining. Even today, many commercial process mining software are developed based on ProM.

ProM is plug-in based open-source process mining software. The plug-ins are applications of algorithms used in process mining. Based on their purpose, plug-ins are divided into five groups: mining, export, import, analysis, and conversions [48]. By executing plug-ins from different groups sequentially, desired process mining analysis are accomplished. For example, basic process mining analysis could use first import plug-in to import XES file, mining plug-in to achieve Petri net model, conversion plug-in to covert Petri net to EPC, analysis plug-in to attain performance graph, and export plug-in to save obtained models. ProM includes plug-ins for all three type of process mining and main perspectives. In year 2016 over 1500 plug-ins had been added to library of proM [15].

When comparing range of possibilities and adaptation, ProM has advantage over other process mining software tools. Unlike commercial process mining tools, which are not easily customizable, ProM was established to support analysis modifying. Open-source plug-in library of ProM allows assembling suitable process mining analysis for different kinds of projects. It has also created the possibility to freely develop new plug-ins to meet needs of new kind of analysis. This has led ProM to be particularly popular among researchers. On the downside, ProM with its customizable nature, is more complicated to use properly than the most commercial tools. This is one reason why in the business sector commercial tools are popular with business process mining.

PM4PY has added a new approach to process mining. Contrary to prior major process mining tools, PM4PY is not a software, but open-source library implemented in Python programming language [47]. It was developed for needs of a truly customizable process mining tool to work on more complex process mining scenarios [49]. Despite being a relative newcomer, PM4PY has already been downloaded over 1 000 000 times and includes all three process mining types and main perspectives [49].

PM4PY can be used similarly to other Python libraries across various platforms, including the command line or terminal, Integrated Development Environments (IDEs), and online platforms. Process mining algorithms in PM4Py are implemented through Python code. The library offers a comprehensive set of tools and functions that enable users to import event log data, execute a range of process mining algorithms, and visualize the results effectively. To fully leverage the capabilities of PM4PY, it requires additional Python libraries such as NumPy, Pandas, Deprecation, and NetworkX [50].

The key advantages of PM4PY relate to seamless integration with Python. As other main process mining tools are primarily design for process mining, Python environment includes a large number of libraries from other disciplines of data analysis such as NumPy, Pandas, and Keras [51, 52, 53]. This allows easy shift from conventional process mining to more complex interdisciplinary analysis including for instance operation recherche and machine learning. Libraries like Pandas also provide superior data editing capabilities compared to software base process mining tools allowing for the entire process mining project to be carried out efficiently on a single platform. For users with prior experience in Python, transitioning to PM4PY for process mining is relatively seamless. In contrast, for individuals without programming experience, PM4PY may present a steeper learning curve compared to other process mining tools.

Commercial process mining tools are designed to provide efficient, user-friendly solutions for process mining. Examples include CELONIS, Microsoft Power Automate, and DISCO, which allow business users to engage in process mining without the need for advanced data science expertise. These tools support rapid analysis and decision-making by automating parts of the data transformation and modeling processes. For instance, DISCO, widely used also in research and academic settings, operates based on the Fuzzy Miner algorithm, producing Fuzzy nets [54]. Its simplicity is enhanced by a single mining model and automated functionality. However, its limited customization capabilities restrict its adaptability for more advanced users seeking tailored approaches. More broadly, commercial process mining tools may show differences in adaptability depending on the type of analysis performed [55]. Consequently, the tools may exhibit varying levels of performance across different analytical tasks.

Thorough and extensive studies comparing the performance differences between various process mining tools are limited. Available research suggests that ProM stands out in terms of methodological depth, making it a robust tool for process mining applications [56, 57]. However, usability of ProM presents a steeper learning curve compared to commercial tools, requiring a higher level of expertise to operate effectively. Latest studies involving PM4PY, indicate that it is at least as efficient as ProM [58, 59]. Furthermore, PM4PY has been rapidly catching up with ProM in terms of feature scope, suggesting it has become a competitive alternative.

2.8 Distinctions Between Process Mining and Data Mining

Process mining and data mining are both integral parts of data science (process mining being also part of process science). Although these disciplines share many commonalities, and process mining is even regarded as having emerged from data mining, each has its own distinct characteristics. The disciplines emerged with a ten-year gap between them, as data mining gained prominence in 1994, and process mining followed in 2004. Despite the modest age difference, data mining has received significantly more attention over the years [60]. While most of the academic world is now somewhat familiar with data mining, the distinctions between process mining and data mining are not commonly known.

Data mining refers to the process of discovering implicit, novel, and potentially beneficial information from vast datasets. [61]. Unlike process mining, which primarily utilizes log data, data mining can be applied to a wide variety of data types. This versatility makes it suitable for a broad range of analyses. Common analyses in data mining seek to uncover association rules, constraints, clusters, and patterns within data that are too complex for humans to detect.

While process mining and data mining share commonalities, such as a data-centric approach, the use of mathematical algorithms, and the goal of providing beneficial yet objective information from data, the perspectives differ. Process mining focuses on fitting and analyzing data within the context of end-to-end process. For this reason, data used in process mining must be well defined and include details about how, where, and when each data point was recorded. In contrast, data mining is less concerned with origin and structure of data and more focused on identifying the broad spectrum of natural patterns it contains. In simplified terms, process mining focuses on both the process and the data it generates, whereas data mining concentrates solely on the data itself. This distinction leads to a difference in perspective between the two disciplines and affects the types of questions that can be addressed using each approach.

Typically process mining aims to address the question of *why*, while data mining focuses more on the question of *what*. However, this distinction is not entirely straightforward. Process mining commonly begins with process discovery, which seeks to answer the question, 'How is this process being performed?' This initial focus is more aligned with the *what* type of question rather than the *why*. As the analysis progresses and incorporates additional types and perspectives of process mining, the focus shifts to seeks to answer the question, 'Why is the process being performed in this manner?' This transition from *what* to *why* is not characteristic of data mining. Data mining employs robust methods for uncovering patterns across a wide range of data types. However, its capability to link these patterns to the underlying processes that generated the data is limited, making it less effective at answering *why* questions.

Process mining and data mining offer distinct perspectives and methodologies. Despite the examples presented earlier, neither discipline is inherently superior. Both offer advanced tools for data analysis and are widely utilized in practice. It is important to recognize the differences between these approaches and select the appropriate tool for the task at hand. In many cases, achieving the best results involves integrating these disciplines and other data analysis techniques.

3 Process Mining in Healthcare

When process mining became more widely used in the early 21st century, healthcare was considered one of its most promising application areas [15]. Although the use of process mining has significantly increased in healthcare research in recent years, its application has not yet become common practice in the field [62]. Much of the research has focused on the development of process mining algorithms and the study of small patient cohorts in specialized healthcare settings [62]. Notably, most of these applications have centered on process discovery or resource assessment [63]. Effective methods for analyzing large scale healthcare data have not yet been widely published. This is partly due to the unique characteristics of healthcare processes, which present distinct challenges that have not yet been fully addressed.

3.1 Distinguishing Characteristics and Challenges

Ten distinguishing characteristics and challenges related to process mining in healthcare have been identified [6]. However, due to partial overlaps with Section 2.5, only the most significant ones relevant to this thesis are discussed in this section.

One primary distinguishing characteristic and challenge in healthcare is the complexity of diverse patient pathways. Unlike many traditional service industries, healthcare does not rely on predefined process pathways. Instead, patients receive personalized care plans based on their specific needs, with healthcare professionals designing these plans while following general guidelines. In the era of personalized medicine, the shift from a generic model to more personalized prediction, prevention, and treatment, has further amplified both the volume and variability of data [64].

This leads to another important challenge: the significance of infrequent behavior. In many industries, infrequent behavior is often ignored as noise. However, in healthcare, due to the complexity of patient pathways and the small margins of error, even infrequent behavior must be considered important. Including infrequent pathways in the analysis can help identify more efficient pathways, as well as incorrect, inefficient, harmful, or even life-threatening ones. While process mining in healthcare generally focuses on larger patient groups, recognizing these infrequent pathways remains essential.

Identifying true infrequent behavior is further complicated by the often poor quality of healthcare data [65]. This is partly due to manual event input, which increases subjectivity and errors. Additionally, a common practice in healthcare, called a workaround, involves recording events later in batches rather than immediately, introducing further timestamp errors [66]. Furthermore, the use of text-based records for some events and the varying recording methods across different devices also contribute to data quality issues.

An important aspect, particularly in the healthcare sector, is that healthcare professionals possess expertise in their field but often lack in-depth knowledge of data science. For process mining results to be effectively applied at the ground level, they must be presented in a manner that is easily understandable without requiring specialized expertise for interpretation. This becomes even more critical when transitioning process mining tools from research to operational use. At the same time, healthcare professionals remain ultimately responsible for patient care. Therefore, the analysis process must be transparent and easily reproducible. This raises questions about newer process mining methods, such as those involving machine learning, which are considered promising in the healthcare field but where the causal relationships behind the results are not always fully understandable.

These and other distinguishing characteristics and challenges presented in [6] must be considered when applying process mining in healthcare settings. Addressing these issues may require the use of custom algorithms or the development of entirely new tools specifically designed for healthcare applications.

4 HUS Helsinki University Hospital

HUS Helsinki University Hospital (HUS) is the biggest health care provider in Finland spanning across multiple facilities in Uusimaa region. It is a major contributor of specialized healthcare in the comprehensive public healthcare system of Finland. HUS holds the responsibility for organizing specialized healthcare and emergency care in the Uusimaa region as well as providing university hospital-level specialized healthcare in Uusimaa, South Karelia, Kymenlaakso, and Päijät-Häme regions. Furthermore, HUS plays a significant role in the nationwide treatment of numerous rare and severe diseases. Approximately 31% of population of Finland (5 603 851 as of 12/2023) resides in the Uusimaa region [67]. This makes HUS responsible for nearly one-third of Finland's specialized healthcare based on population.

4.1 Pediatric and Adolescent Emergency Departments

In Finland, tertiary-level emergency care, pediatric in-hospital care, and pediatric intensive care are exclusively provided in public hospitals. HUS has centralized its pediatric and adolescent emergency departments to two locations in the Uusimaa region: New Children's Hospital (NCH) in Helsinki and Jorvi Hospital in Espoo. These emergency departments serve 315 250 residents (as of 12/2023) aged 0-16 years in the Uusimaa region and provide comprehensively specialized healthcare [67]. In both emergency departments, pediatric, psychiatric, surgical, and neurological specialties are served. In addition, NCH has the most extensive pediatric sub-specialty readiness in Finland. Most challenging and acute cases from Uusimaa and, in some cases, from all over Finland are directed here for emergency care.

New Children's Hospital and Jorvi Hospital serve around-the-clock specialized emergency care every day of the year. This is intended for acute life-threatening situations and serious urgent health-threatening conditions. In addition, both units also provide pediatric and adolescent primary healthcare emergency services outside of office hours from 4 pm to 10 pm during weekdays. During office hours from 8 am to 4 pm, pediatric and adolescent primary healthcare emergency services are provided alongside the emergency services at local health centers. On weekends and holidays, primary care emergency services are centralized to larger hospital units in Uusimaa region. During these days, pediatric primary care emergency services are provided at both the NCH and Jorvi Hospital from 8 am to 10 pm. These primary care emergency services are intended for acute conditions requiring rapid treatment that do not however require specialized emergency care.

Pediatric and adolescent emergency departments at NCH and Jorvi Hospital operate with varying staffing levels and bed capacities depending on the time of day. At NCH, an average of 18 inpatient beds are typically available, along with five dedicated overnight observation beds. During office hours, the ED is staffed by two senior physicians and five resident physicians. During on-call hours, staffing levels vary. However, the standard includes two primary care physicians and six specialists from different fields. As a university hospital, multiple backup physicians from multiple specialty fields are also available. NCH also hosts the only pediatric and adolescent intensive care unit in the HUS region, receiving all critically ill patients requiring intensive care. At Jorvi, the facility includes seven inpatient beds and three designated overnight observation beds. The ED is consistently staffed by a team of three physicians, comprising one senior physician and two residents, one of whom is assigned to primary care during on-call hours.

4.2 Triage

In emergency department context, triage refers to a systematic process aimed at promptly evaluating the severity of illness or injury of a patient, assigning priorities, and directing each patient to the appropriate treatment area [8]. Historically triage system rose from military medicine as a method to rapidly identify and allocate wounded soldiers to groups of beyond saving, needs urgent care, and can wait [68]. From there, the triage system was adapted to emergency departments to match limited care resources to a growing number of patients, of whom only a portion require immediate care.

In real life settings, every patient cannot get immediate treatment in emergency departments due to limited care resources. Furthermore, prioritizing patients based solely on their order of arrival is not sensible given the varying severity of illnesses and injuries. For this purpose, triage systems were implemented for continuous process to establish a structured order in which patients should be treated in emergency departments. Although there are differences in various triage systems, the basic principle is that when a patient arrives in the emergency department, the triage nurse defines the triage level for the patient. This level determines treatment urgency and at what stage, compared to other patients, the patient should be treated. The aim is to ensure that the most critically ill and injured patients receive urgent treatment, even if care resources are in full use.

4.2.1 ESI – Emergency Severity Index

NCH and Jorvi hospital use Emergency Severity Index (ESI) triage system to manage patient flow at pediatric and adolescent emergency departments. ESI, developed in the United States in the late 1990s, is among several 5-level triage systems that classify patients based on urgency and anticipated resource requirements [69, 70]. The ESI scale ranges from 1 to 5, with level 1 indicating the most urgent cases and level 5 the least. ESI levels 1 and 2 are determined solely based on the urgency and clinical condition of the patient, while ESI levels 3, 4, and 5 are assigned according to predicted resource requirements [69].

ESI level 1 is designated to patients presenting to the ED in critical condition who require immediate, life-saving interventions. Such interventions include airway management, respiratory support, emergency medications, and hemodynamic support, such as fluid resuscitation or blood transfusions [69]. Clinical scenarios that necessitate lifesaving interventions encompass cases where patients are intubated, unresponsive, pulseless, apneic, in severe respiratory distress, or experiencing profound hypotension or hypoglycemia [69]. ESI level 2 is designated to patients who are not in immediate life-threatening condition but are at high risk of quickly progressing towards it. These patients include those who have experienced a high-risk situation, or are confused, lethargic, or disoriented, or are experiencing severe pain or distress [69]. Age adjusted abnormal values for heart rate, blood pressure, oxygen saturation, and pediatric fever also meet the criteria for ESI Level 2.

ESI 3, 4, and 5 levels are assigned based on resource requirement estimation [69]. These patients are generally stable enough to wait for treatment with low risks for a rapid deterioration in condition. In the assessment process, the triage nurse evaluates and estimates the number of resources the patient is likely to require. ESI level 3 patients are estimated to need two or more resources, while ESI level 4 patients are estimated to require one resource, and ESI level 5 patients are estimated to need no resources. The categorization of emergency department resources considered in ESI scoring, as well as those that are not, is detailed in Table 5.

Resources are assessed based on the type of resource needed rather than individual tests [69]. For example, all tests performed on blood and urine are considered a single laboratory test resource, regardless of the number of tests conducted. Similarly, different types of imaging conducted using the same imaging modality are counted as a single resource. However, imaging procedures performed using different modalities each constitute a separate resource.

Table 5: ESI resources. An example of ESI resource categorization (Based on [69])

Resources	**Not resources**
Laboratory tests (blood, urine)	History check and physical examination
ECG, X-ray, CT, MRI, Ultrasound	Point-of-care testing
IV fluids (hydration)	Saline or Heplock
IV, IM or nebulized medications	PO medications, Prescription refills
Speciality consultation	Phone call to primary care physician
Simple procedure (= 1 resources)	Simple wound care
Complex procedure (= 2 resources)	Crutches, Splints

Although ESI levels 3, 4, and 5 are determined based on predicted resource requirements, the ESI system is not a measure of workload. Instead, the predicted resource requirements are used in ESI as indicators of patient acuity. Clinical studies on this topic have shown a correlation between resource utilization and patient acuity [8]. Additionally, the ESI system is not intended for monitoring patient condition but serves solely as an initial assessment of patient status in the ED [69]. Therefore, the ESI score should not be changed once treatment has begun. Table 6 includes example pediatric trauma cases for all five ESI levels.

Table 6: Trauma ESI examples. Example pediatric ED trauma cases for all ESI levels (Based partly on [69]).

Patient	**Resources**	**ESI**	**Reasoning**
EMS arrives with an 8-year-old child who was hit by car while biking. Patient is unconscious, pale, and has labored respiration.	Lifesaving intervention, ESI assessment without considering resources.	ESI 1	Life threatening trauma.
EMS arrives with a 15-year-old adolescent who hit his head hard on the edge of the rink during a hockey match. The patient is awake, has stable basic vital signs, and demonstrates motion in all limbs. Currently patient is immobilized.	High risk injury, ESI assessment without considering resources.	ESI 2	Injury mechanism suggest a high-risk injury.
EMS arrives with a 12-year-old child who tumbled on the trampoline and fell on left hand. The left hand shows an obvious deformity between the elbow and wrist. Radial and ulnar pulses are normal and fingers are warm on the left hand. Basic vital signs are normal and pain is 5/10.	Need more than one resources, low risks for a rapid deterioration in condition.	ESI 3	Fracture will require reduction. Also X-ray, Lab test, IV antibiotics, and pain medication are needed.
10-year-old child arrives ED brought by his father. The patient fell during play and sustained a 2 cm superficial laceration to his knee. Basic vital signs are normal.	Need one resources.	ESI 4	Laceration will need suturing.
Scared mother brings 3-year-old child ED who fell from the sofa. No signs of injury. Basic vital signs are normal and the child seems playful.	No need for resources	ESI 5	Physical examination and reassurance of the mother are required.

5 Literature Review

In recent years, process mining has been applied across a range of healthcare applications, which multiple systematic literature reviews have explored extensively [71, 72, 62, 73, 74]. However, applying process mining in emergency department settings remains particularly challenging due to the high variability of patient cases. This complexity, even relative to other healthcare domains, has resulted in limited research on its application in ED settings. Thus, research on process mining in pediatric and adolescent ED is even more limited. In this chapter, previous research on process mining in pediatric and adolescent ED is reviewed, along with studies focused on ED settings more broadly.

5.1 Key Research of Process Mining in Pediatric ED

Study conducted by A. B. Durojaiye et al. [75] investigates the use of process mining to analyze the in-hospital flow of pediatric trauma patients across multiple care locations. Utilizing a cohort of n = 1 941 pediatric trauma patients from a Level I trauma center with the two highest activation levels, Alpha and Bravo, this study maps patient pathways and transitions across care settings. The Flexible Heuristics Miner algorithm was used to generate process models, uncovering 28 distinct patient pathways and 20 primary care transitions, highlighting potential areas for optimizing patient flow.

Another study conducted by A. B. Durojaiye et al. [76] investigates collaborative practices in pediatric trauma care. The study method combines process mining from an organizational perspective with network analysis. It uses the igraph package in R for analysis. The findings suggest that closer collaboration among healthcare providers in the pediatric ED reduces ED length of stay (LOS).

Study conducted by R. C. Basole et al. [77] investigates how different treatment approaches affect outcomes for pediatric asthma patients. The researchers created an interactive algorithm for visual analytics focused on process exploration and discovery. Algorithm was utilized with Gephi tool to examine clinical data from n = 5 784 pediatric asthma patients treated in the pediatric ED. The study emphasizes the complexity of pediatric asthma care processes. More broadly, it highlights the potential of visual analytics to enhance understanding and drive quality improvement efforts within ED settings.

5.2 Key Research of Process Mining in ED

Study conducted by D. Duma et al. [78] proposes a new framework for process mining tailored to map and analyze ED workflows using customized process discovery tools. The framework aims to produce a simple yet accurate model that reflects the diverse patient paths in the ED. This approach is tested with n = 88 272 ED patient data, providing both a retrospective analysis and comparisons with standard process mining methods. The study emphasizes the benefits of customized process mining for improving ED efficiency and optimizing resource allocation.

Study conducted by F. Rismanchian et al. [79] combines process mining with optimization to improve ED layout by reducing unnecessary patient travel distances. Using process mining software DISCO to analyze clinical data, it identifies inefficient spatial assignments and then applies a genetic programming (GP) approach to reconfigure unit placements. Study shows that optimizing ED layout can significantly reduce patient travel distances, offering potential to enhance efficiency in both existing and newly designed EDs.

Study conducted by R. Andrews et al. [80] focuses on improving ED patient flow, addressing challenges like overcrowding, prolonged LOS, and access block, which negatively impact patient outcomes. Using process mining software DISCO with the BPM Lifecycle framework, the authors analyzed n = 1 473 chest pain cases to identify factors affecting patient flow. The study identified only minor differences between short LOS and prolonged LOS patients through process mining, with waiting times for hospital admission from ED playing a significant role in increasing LOS. Essential methodological observations emphasize the need for better data quality and the importance of collecting start and end times for accurate performance analysis.

Study conducted by C. Alvarez et al. [81] utilizes DISCO process mining software to examine role interactions among ED professionals from an organizational perspective. The methodology was applied to a dataset of n = 7,160 ED cases sourced from hospital information system (HIS). The study indicates that ED professionals increase their level of collaboration as patient severity rises. Although the study results are somewhat predictable, the advantage of the proposed methodology is its ease of replication for other study questions.

Study conducted by M. Cho et al. [82] proposes a framework for assessing ED processes using process mining, incorporating 16 performance indicators (EDPPIs) based on the four key perspectives of the devil's quadrangle: time, cost, quality, and flexibility. The methodology was validated using a dataset of approximately 30 000 ED patients, leveraging process mining tools such as DISCO, ProDiscovery, and ProM. The study underscores the practical value of using process mining tools to enhance the efficiency of ED operations.

Study conducted by T. G. Erdogan et al. [83] applies multi-perspective process mining techniques to evaluate emergency processes in a university hospital, based on a dataset of n = 894 ED cases. The study employs the Goal-Question-Feature-Indicator (GQFI) method with process mining tools DISCO and R-bupaR to assess time sequences within ED processes. Through this analysis, the study identified deviations such as skipped triage and consultation request steps, as well as two key bottlenecks in the emergency process. The study demonstrates that multi-perspective process mining can be an effective approach for identifying inefficiencies in ED processes.

Study conducted by E. Rojas et al. [84] utilizes process mining to analyze ED case performance, focusing on identifying factors that contribute to LOS. Study was conducted using DISCO and used data of n = 7 160 ED cases. The study indicated that a key driver of increased LOS is when patients enter a recurring loop, alternating between examination and treatment steps.

Study conducted by F. Davari [85] compares process mining and simulation to identify bottlenecks and optimize patient flow in the ED. A dataset of n = 1 275 ED cases, selected from a population of 39 264 using systematic random sampling, was analyzed using ARENA simulation software and the DISCO software. The results highlighted key inefficiencies, including delays in processing orders from doctors, waiting for test results, and discharge congestion. The findings from this study underscore the utility of simulation techniques and process mining for supporting data-driven, resource-aligned decision-making in ED optimization.

Study conducted by G. Ibanez-Sanchez et al. [86] investigates the use of process mining to support value-based healthcare by analyzing emergency processes, specifically in stroke cases. The study utilized a dataset of n = 9 046 emergency cases from 2 145 stroke patients applying Process Mining with PMApp tool and a Question Driven methodology to understand treatment flow. The results demonstrate that process mining effectively highlights differences in stroke patient flow compared to other emergency cases, identifying critical timing factors. More broadly, the study highlights the benefits of process mining in identifying variations within ED patient processes.

5.3 Novelty of the Study and Tabular Overview of Reviewed Articles

Research on pediatric and adolescent emergency departments using process mining is highly limited, highlighting the inherent value of this study. More broadly, no prior research has applied process mining to an emergency department with a comparable patient volume. Existing studies predominantly focus on specific patient cohorts rather than comprehensive datasets. Among the only two studies conducted with larger patient volumes, the focus was either on algorithm development or on the use of external efficiency metrics beyond process mining [78, 82]. This study, however, investigates how current process mining methods and tools can be effectively applied to large scale data. It provides novel insights into pediatric and adolescent emergency department operations, process mining with large patient volumes, and the strengths and limitations of current methodologies.

To provide a clearer overview, a tabular summary is presented to encapsulate the research discussed in this chapter. The summary, shown in Table 7, includes the medical field of the study, the focus of the research, and the miners and tools utilized.

Table 7: Summary of literature review.

Study	Medical Field	Study Focus	Miner and Tool
[75]	Pediatric ED, Pediatric Trauma Cohort	Patient Pathway	Flexible Heuristics Miner (FHM), ProM
[76]	Pediatric ED, Pediatric trauma Cohort	LOS, ED Professional Collaboration	igraph, R
[77]	Pediatric ED, Pediatric Asthma Cohort	Visual Analytics	Custom Algorithm, Gephi
[78]	ED	Process Flow, Prediction	Heuristic Miner (HM), Inductive Miner – infrequent (IMi), ProM
[79]	ED	Layout optimization	Fuzzy Miner, DISCO
[80]	ED, Chest Pain Cohort	Process Flow, LOS	Fuzzy Miner, DISCO
[81]	ED	ED Professional Collaboration	Fuzzy Miner, DISCO
[82]	ED	Process Performance Indicators	Fuzzy Miner, DISCO, ProDiscovery, ProM
[83]	ED	Process Improvement	Fuzzy Miner, DISCO, R-bupaR
[84]	ED	LOS	Fuzzy Miner, DISCO
[85]	ED	Process flow	Fuzzy Miner, DISCO
[86]	ED, Stroke Cohort	Value-Based Healthcare	Custom Algorithm, PMApp

6 Methodology

This study consists of four methodological components: data filtering, event log structuring, process mining, and analysis. The primary objective is to investigate how process mining can be applied to generate patient pathway models using existing miners for process discovery in large datasets from pediatric and adolescent ED. Additionally, the study aims to investigate resource allocation and systematically identify failure demand within the ESI triage classification. The ED resources considered in this study include procedures, imaging, and laboratory tests. Medical checks performed by physicians or nurses, for which no procedure is recorded, are excluded due to insufficiently clear data for generating accurate models. To achieve these objectives, customized Python algorithms will be developed for data preprocessing. Four widely used miners will be tested using ProM for process discovery. The most effective miner will then be selected to conduct process discovery on pediatric and adolescent ED data, categorized by ESI triage levels. The results will be analyzed to provide insights into ED operations.

6.1 Data

HUS patient data collected from the Hospital Information Systems (HIS) has been utilized in this study. The dataset includes visits and medical records for individuals aged 0 to 16 years at HUS units, spanning from January 1, 2021, to December 31, 2023. The data was extracted from the HUS Azure Data Lake, containing electronic health records from Apotti and laboratory test data from Multilab. The extracted data includes multiple datasets, with the most essential for this study being demographics, visits, ESI triage scores, procedures, imaging, and laboratory tests. Of these, visits, ESI triage scores, procedures, imaging, and laboratory tests include timestamps marking the moments of execution, which are vital for event log formatting. The procedure dataset is the most heterogeneous, encompassing a wide range of data, from minor procedures and diagnostic actions, such as catheter insertion and ECG tests, to major surgeries and other high complexity interventions involving multiple personnel. Additionally, all datasets contain numerous other attributes that are utilized for statistical analysis.

To apply discovery mining techniques, pediatric and adolescent ED data must first be isolated from the dataset and structured as an event log. Given the inherent noise and incompleteness of healthcare data, achieving perfectly accurate filtering and event log formatting at this dataset scale is not feasible [6, 87]. For this study, the pre-processing phase proved to be the most challenging and time consuming task. The best results were achieved through trial and error and close collaboration with physicians working in the ED.

Due to the highly sensitive nature of healthcare data, careful consideration must also be given to information security measures [88]. The data utilized in this study is governed by the Finnish Secondary Use Act. Therefore, obtaining HUS study permission was a necessary task and a non-disclosure and information security agreement needed to be signed. Research permission was granted in June 2023. The data needed to be pseudonymized and managed on a secure platform. The whole data processing was conducted on secure HUS Acamedic virtual machine platform. HUS Acamedic is a secure virtual operating environment that complies with the Finnish Secondary Use Act. Only upon completion of the study were results that complied with FINDATA (The Finnish Social and Healthcare Data Permit Authority) regulations exported from the platform.

6.1.1 Filtering

The initial filtering of data was carried out in four parts, as shown in Figure 4. After filtering, n = 205 083 pediatric and adolescent ED patient cases were obtained for analysis. The filtering was performed using Python 3.8.10 and the Pandas library (version 2.0.3). The following section aims to clarify the purpose of certain filtering criteria and the possible trade-offs that could arise from implementing them.

The requirement for an ESI triage classification (or that the patient visit type be designated as Referred) is driven by the aim of selecting only emergency patients for this study. The Pediatric and Adolescent ED at the New Children's Hospital and Jorvi Hospital also treats a limited number of non-emergency patients, such as certain follow-up or additional visits, which this analysis aims to exclude. ESI classification is not assigned to these patients in the emergency department. However, the filter also removes any potential emergency visits where the ESI classification was not recorded due to either human or technical reasons.

To ensure that only those tests, imaging and procedures related to the ED visits are included in the study, they have to be marked as performed before the ED visit is marked as finished. Tests or procedures ordered from ED as part of follow up care for patients transferred to inpatient care are not to be included in the study. However, in a fast rapidly evolving environment like ED, significant delays may occur between documentation and the actual performance of certain actions. As a result, this filtering process could potentially lead to the omission of some ED related laboratory tests, imaging, and procedures.

The requirements that laboratory tests, imaging, and procedures are recorded as ordered specifically by the Pediatric and Adolescent ED at either NCH or Jorvi Hospital, ensures that actions taken prior to ED admission but sharing the same hospitalization episode code are excluded from the dataset. These situations may arise when patients are transferred from another unit or, in some cases, from within the same hospital to the ED. Since the focus of this study is solely on ED processes, actions occurring prior

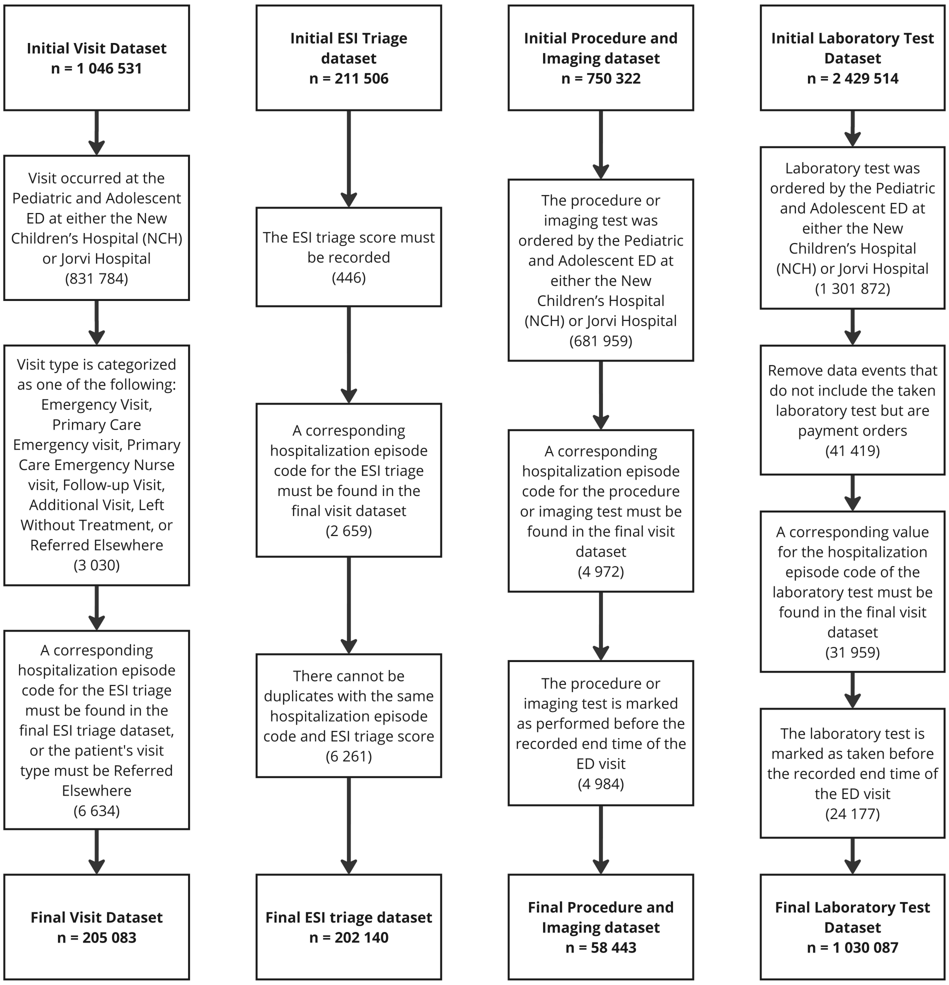

Figure 4: Filters used to preprocess datasets. Number of events removed by the filter are shown in parentheses (n).

to the ED admission are intentionally omitted from the data. An alternative, more intuitive approach could be to require that laboratory tests, imaging, and procedures are recorded as performed only after the ED visit is marked as started. However, especially in acute situations, it might occur that actions are performed and documented in real time for the patient before the official start of the ED visit is recorded. To avoid distorting process data for acute patients, the requirement based on the orderer was applied for filtering. The trade-off in this approach is that any actions ordered from other hospital units such as after a consultation and performed in the ED will be excluded from the dataset.

The requirement for ESI triage, ensuring there are no duplicates with the same hospitalization episode code and ESI triage score, is intended to remove unnecessary ESI recordings that are irrelevant for this study. In certain situations, in the HUS pediatric and adolescent ED, a new ESI triage entry is logged for a patient, even though the ESI triage score remains unchanged. Such situations occur, for instance, when a patient transitions between specialties in the ED, such as from pediatrics to surgery. Another example is when the ED shifts from regular office hours to on-call hours, which can also trigger a new triage entry without changing the ESI triage score. Since these entries are more related to the operational methods of the HIS and fall outside the scope of this study, in such situations only the first ESI triage entry is included in the dataset.

In situations where ESI triage score of a patient changes during the ED visit (n = 5270), all corresponding entries are included in the dataset. This results in cases within the ESI triage dataset where multiple data points share the same hospitalization episode code but have different ESI triage scores. In event log, and consequently in process mining, this situation is represented as a change in the ESI triage score. However, in statistical analysis and clustering, each ED visit must not have more than one ESI triage score to prevent duplication in the statistics and clusters. As stated in section 4.2.1, ESI triage score should not be changed after treatment process has been initiated. Therefore, for statistical analysis and clustering, ESI triage filtering is applied to patients who experience a change in their ESI score during the ED visit as follows: If the score change occurs within 15 minutes of the initial ESI entry, the later score is retained. However, if the score change occurs after 15 minutes, the first ESI score is retained. This approach enables the quantification of how frequently different initial ESI triage scores are assigned during ED visits and how often score changes occur across the various ESI categories. Since a late change in the ESI score indicates either a significant error in assessing resource demands, significant misjudged immediate care needs, or a clear change in the condition of patient during the ED visit, this group of visits holds significant relevance in future study.

6.1.2 Event Log Structuring

The event log structuring was conducted as illustrated in Figure 5, resulting in event log containing n = 467 621 events. This process utilized Python 3.8.10 alongside the Pandas library (version 2.0.3). The following section outlines the steps taken during structuring and the rationale behind each step.

The event log structuring steps for the datasets of procedures, imaging, and laboratory tests are designed to account for the fact that a single action in the ED can result in multiple entries in the HIS. For example, a single blood test led to several entries when multiple analyses are conducted for sample. This is captured in the dataset as multiple blood tests recorded with identical or closely matching timestamps. Timestamp delays are often longer when different types of tests, such as point-of-care (POC) tests, blood tests, or urine tests, are performed simultaneously. To accurately represent the care pathway, the focus should be on capturing the actual care processes rather than strictly adhering to the sequence of recorded events. Therefore, each care action must be consolidated into a single, unified entry in the event log.

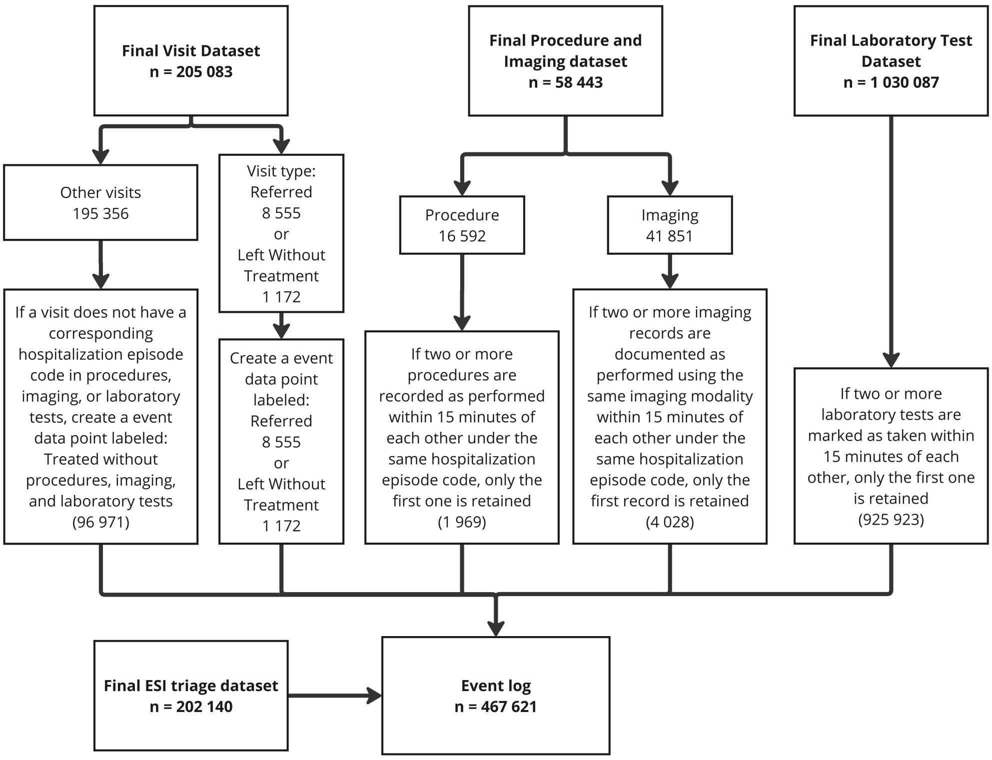

Figure 5: Structuring of event log for process mining analysis. Number of events removed by the condition are shown in parentheses (n).

In this study, a threshold of 15 minutes has been set, within which multiple procedures, multiple imaging actions using the same imaging modality, and multiple laboratory tests are recorded as a single event. The aim is to eliminate the above-mentioned discrepancy between the data points and the actual care actions. The threshold is based on data exploration, trial and error, and collaboration with clinicians. Although this threshold is critical for effectively and meaningfully applying discovery mining techniques, it may introduce trade-offs.

ESI triage, laboratory tests, procedures, and imaging represent the main events recorded during pediatric and adolescent ED visits. However, approximately half of these visits include no entries beyond the initial ESI triage and the ED visit record itself. To address this, three additional artificial event types were introduced from visit data: Treated without procedures, imaging, and laboratory tests, Left without treatment, and Referred. These event types serve to distinguish this patient group and facilitate a clearer interpretation of the process mining models.

The final step is to combine all datasets into one unified event log. In this combined dataset, an activity attribute is created for each event, named according to the event class as follows: ESI Triage, Laboratory Test, Procedure, and Imaging. Additionally, three previously defined artificially created event activity groups are included. The case ID is based on the hospitalization episode code, and the timestamp is the moment each event was recorded to HIS. For the artificially created actions, the timestamp is set to the end time of the visit.

6.1.3 Statistics

In this section, basic statistics are presented from the filtered datasets to provide insight into the data and support the analysis of the process model results. Table 8 summarizes some key data metrics. In the table, ICU refers to the Intensive Care Unit, HDU to the High Dependency Unit, and OR to the Operating Room. Figure 6 illustrates four time series of pediatric and adolescent emergency department visits. These time series display visits by hour on weekdays, by hour on weekends, by day of the week, and by month. The time series differentiate between various ESI score levels to show how patient volumes are distributed across different ESI levels.

Table 8: Basic statistics of filtered datasets.

Sex	**n**	**%**
Male	112 455	54.8
Female	92 605	45.2
Not available	23	≈ 0
Age (Years)	**n**	**%**
0 – 3	93 213	45.5
4 – 7	43 357	21.1
8 – 11	32 988	16.1
12 – 16	35 525	17.3
Native language	**n**	**%**
Finnish or Swedish	148 848	72.6
Other language	55 944	27.3
Not available	291	0.01
Place of residence	**n**	**%**
Metropolitan area	179 421	87.5
Other HUS regions	17 874	8.7
Rest of Finland	4 905	2.4
Abroad	2 883	1.4

ESI triage	**n**	**%**
ESI 5	69 505	33.9
ESI 4	71 903	35.1
ESI 3	40 506	19.8
ESI 2	13 867	6.8
ESI 1	1 001	0.5
ESI not assigned	8 301	4.0
Top visit causes	**n**	**%**
Fever	35 064	17.1
Hand symptoms	10 658	5.2
Wound	9 388	4.6
Abdominal pain	9 332	4.6
ED disposition	**n**	**%**
Discharge	191 006	93.1
Ward	12 901	6.3
ICU	280	0.4
HDU	124	0.1
OR	772	0.1

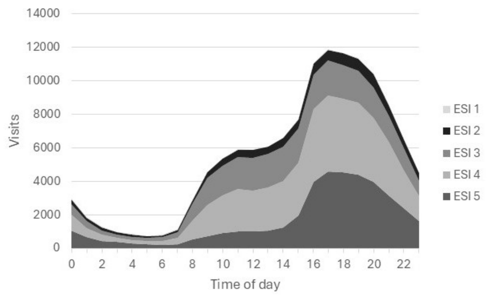

(a) Weekday visits by hour.

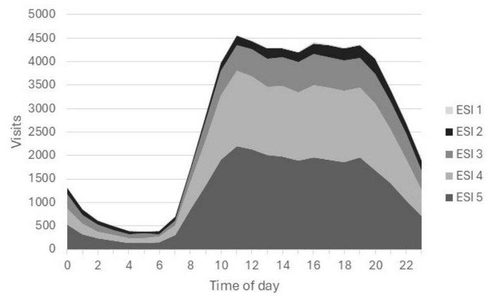

(b) Weekend visits by hour.

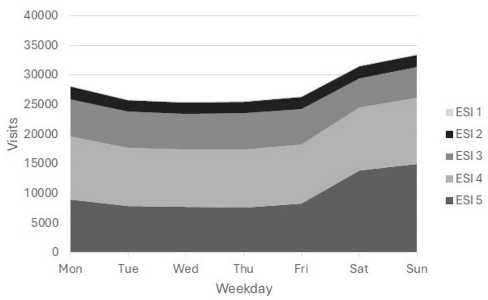

(c) Visits by day of the week.

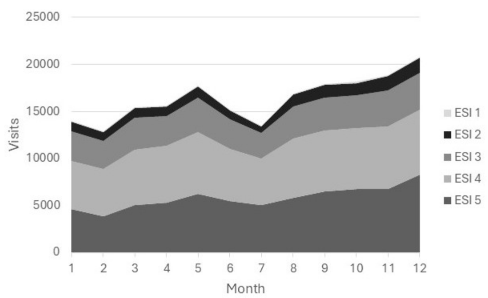

(d) Visits by month.

Figure 6: The time series of pediatric and adolescent ED visits, with the ESI triage categories separated.

6.2 Process discovery

Process discovery in this study was conducted using ProM 6.12. It was chosen for its versatile methodological capabilities and adaptability to various types of analyses. A comparative analysis with PM4PY was not performed due to data protection restrictions that prevented the installation of required libraries on the HUS academic platform. Various process discovery algorithms were tested on the event log to identify the miner most suitable for this study. The tested algorithm plug-ins included Mine for Heuristic Net using the Heuristic Miner (Heuristic Miner), Interactive Data-Aware Heuristic Miner (iDHM), Mine for Fuzzy Model (Fuzzy Miner), and Convert Log to Directly-Follows Graph (DFM Miner). Before applying each miner, the Add Artificial Event plug-in was used to insert artificial start and end events for each case in the event log to enhance the clarity of the mined models.

Each miner was first applied to the event log using the default settings. Heuristic Miner, iDHM, and Fuzzy Miner were further tested with various customization configurations. In contrast, DFM was used without any filtering options, reflecting the process exactly as sequenced in the event log. The resulting process models with the default settings for all tested miners are shown in Figure 7. As expected, Heuristic Miner, iDHM, and Fuzzy Miner face challenges in generating process models that are both sufficiently precise and clear from the event log. These models do not capture short- or long-term dependencies between actions in a clearly interpretable manner. In contrast, DFM clearly visualizes short-term dependencies between actions. However, identifying long term dependencies between actions remains profoundly challenging in the resulting model.

A major reason traditional miners struggle with the event log in this study is the lack of a clear directional flow between actions. The event log shows a considerable frequency of trace occurrences in both directions between actions, making it difficult to identify significant correlations and causal relationships. As a result, even after careful tuning, Heuristic Miner, iDHM, and Fuzzy Miner, do not produce meaningful results. DFM Miner stands out as the only one capable of managing the complex immediate relations between actions, despite utilizing the simplest algorithm. For this reason, the DFM Miner was chosen for the final study. However, due to its limitations in capturing long term dependencies and representing complete processes, the Filter Book plug-in in ProM is utilized with it. This plug in sequences all cases in the event log and displays the absolute frequency of all case sequences.

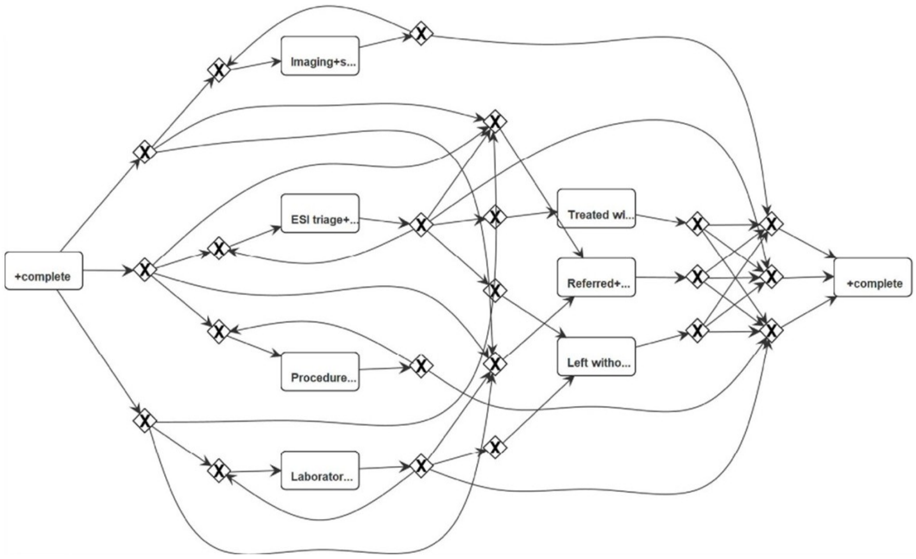

(a) Heuristic Miner.

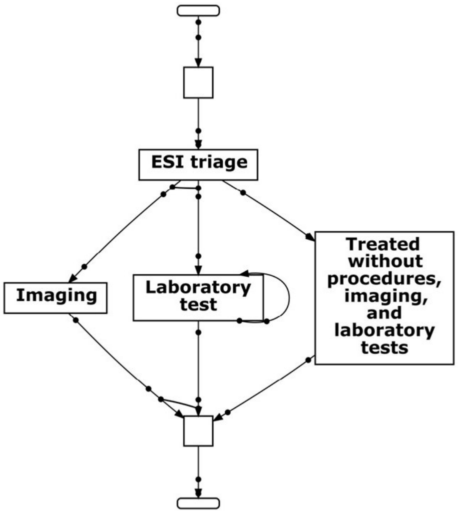

(b) iDH Miner.

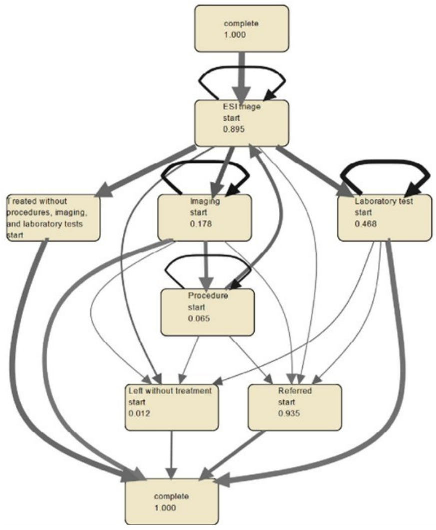

(c) Fuzzy Miner.

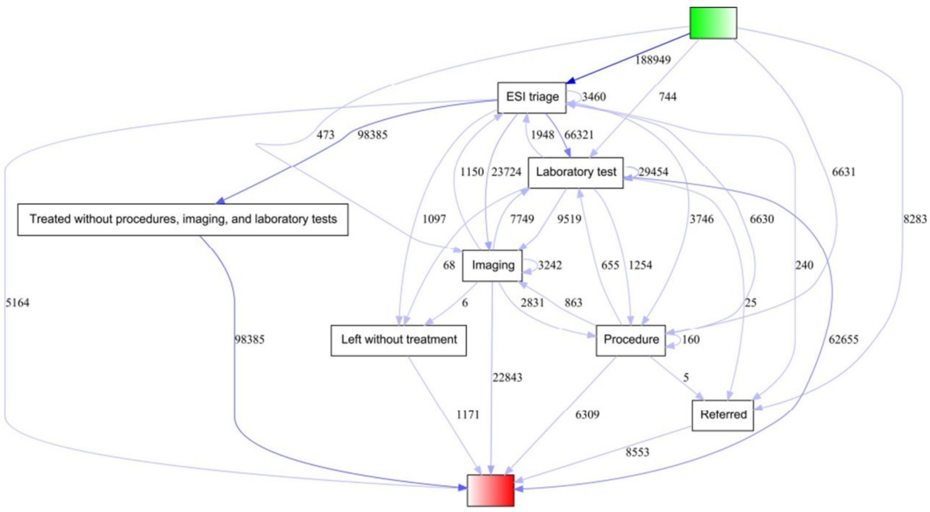

(d) DFM Miner.

Figure 7: Process models generated by four process discovery miners with default settings.

7 Results

The DFM Miner was used to generate process models for the entire event log as well as for each ESI score cluster. These models are presented in Figures 8, 10, 12, 14, 16, and 18. Following FINDATA regulations, all connections showing fewer than 5 occurrences have been removed from the models. In total, this led to the removal of 37 traces across all process models. Upon initial examination, the process models display a complex trace structure. However, a systematic analysis of the different components of the models yields valuable insights.

When initially examining patient pathways that do not include recorded procedures, imaging, or laboratory tests, a significant difference is observed across the ESI clusters. The least urgent ESI categories, ESI 5 and ESI 4, are to a great extent characterized by traces without recorded procedures, imaging, or laboratory tests, representing 72 % of all ESI 5 visits and 46 % of ESI 4 visits. In contrast, for ESI 3 and ESI 2, this proportion decreases to 29 %. For ESI 1, it is as low as 9 %. Since pediatric and adolescent ED visits are more concentrated in the higher (less urgent) ESI triage score categories, it follows that 48 % of all visits are managed without recorded procedures, imaging, or laboratory tests.

Next, during the examination of patient pathways involving procedures, imaging, and laboratory tests, it becomes evident that some transitions occur more frequently than others. This is illustrated by Figures 9, 11, 13, 15, 17, and 19, which show the four most common pathways for the entire event log as well as for each ESI score clusters. In these figures, ESI refers to the triage assignment, TWPIL represents treatment without procedures, imaging, or laboratory tests, LAB stands for laboratory tests, and IMG stands for imaging. These figures show that the four most common pathways remain consistent across the ESI 5, 4, 3, and 2 clusters, with only variations in order and frequency. Contrary to the initial impression of a complex trace structure that the process models convey, these pathways account for a significant portion of all visits: 93 % in the ESI 5 cluster, 84 % in ESI 4, 66 % in ESI 3, and 62 % in ESI 2. Together, these four pathways account for 78 % of all visits.

The ESI 1 cluster, however, represents a notable exception to this trend. The four most frequent patient pathways, shown in Figures 19, differ in part from other ESI clusters and constitute only 32 % of the ESI 1 visits. The differences in patient pathways for ESI 1, compared to other ESI clusters, are even more pronounced when examining the variety of patient pathway variants within each cluster. With 298 distinct variants, the ESI 1 cluster displays a ratio of 0.30 variants relative to the total number of visits. This means there are 30 distinct patient pathways per 100 patients. In comparison, the ratios are 0.0034 for ESI 5, 0.0078 for ESI 4, 0.017 for ESI 3, and 0.033 for ESI 2. Although the ratios are not directly comparable due to the significant difference in visit numbers between the ESI clusters, this highlights the more individualized nature of patient pathways for ESI 1 patients. Across all visits, the ratio of patient pathway variants to the total number of visits is 0.0063.

Another significant metric that can be derived from process models is the proportion of procedure, imaging, and laboratory test events relative to the number of visits. This metric does not indicate absolute resource consumption but rather reflects the utilization of care and diagnostic resources during ED visits. For instance, two patients who both undergo identical blood test analyses consume the same absolute amount of laboratory resources. However, if the tests for one patient are conducted in a single session and for the other in two separate sessions, the utilization of care resources differs. This proportion of procedure, imaging, and laboratory test events relative to the number of visits demonstrates an inverse relationship with the decreasing ESI score scale. However, the progression across ESI clusters does not follow a linear trend. In the ESI 5 cluster, the proportion of procedure, imaging, and laboratory test events per visit is 0.34, increasing to 0.77 for ESI 4, 1.3 for ESI 3, 1.5 for ESI 2, and reaching 3.1 for ESI 1. A clear increase is observed between the ESI clusters from ESI 5 to ESI 4, from ESI 4 to ESI 3, and from ESI 2 to ESI 1. However, the ratio changes only slightly between the ESI 3 and ESI 2 clusters. The proportion of procedure, imaging, and laboratory test events per visit for all visits is 0.77.

To facilitate further assessment of differences between ESI clusters, the length of stay (LOS) has also been calculated for all visits and each ESI cluster. The most significant increase in LOS is observed between ESI 4 and ESI 3, with a difference of nearly one and a half hours (86 minutes). The differences between ESI 5 and ESI 4, as well as between ESI 3 and ESI 2, are less than one hour (52 and 43 minutes, respectively). The difference between ESI 2 and ESI 1 is only 6 minutes. LOS for all visits is 138 minutes. In this study, the LOS is measured from the point when the ED is recorded as having started (often during triage) to the point when the ED is recorded as having ended, thereby reflecting the patient's treatment time in the ED. It is important to note that the time a patient spends in the ED may be considerably longer, particularly for non-urgent cases, due to potential delays between the arrival and the triage assessment. Tables 9, 10, 11, 12, 13, and 14 present the summary statistics provided in this section for all visits and each ESI cluster.

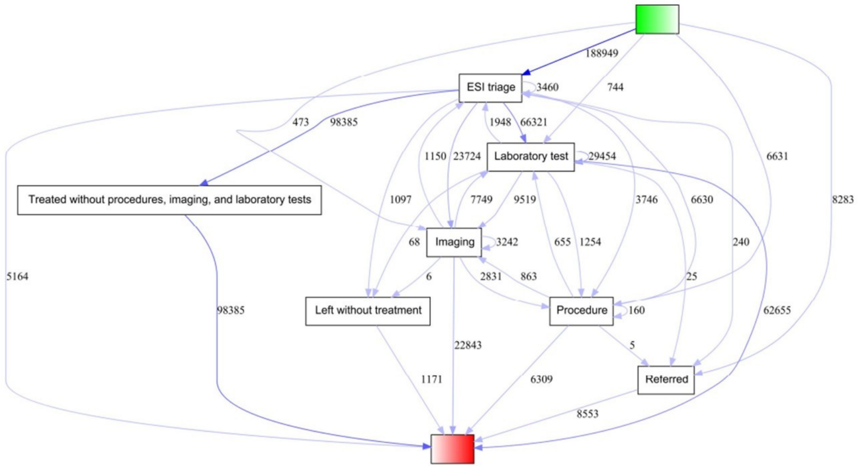

Figure 8: Process model for all visits.

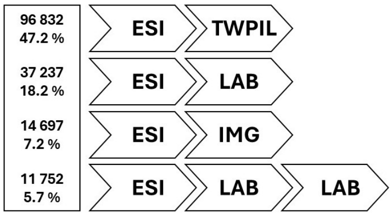

Figure 9: Top 4 pathways for all visits.

Visits	205 083
LOS (min)	139
Pro, Img, Lab events per visit	0.77
Trace variants	1 301
Visits of top 4 traces	160 518
Visits of top 4 traces (%)	78.3

Table 9: Key Metrics for all visits.

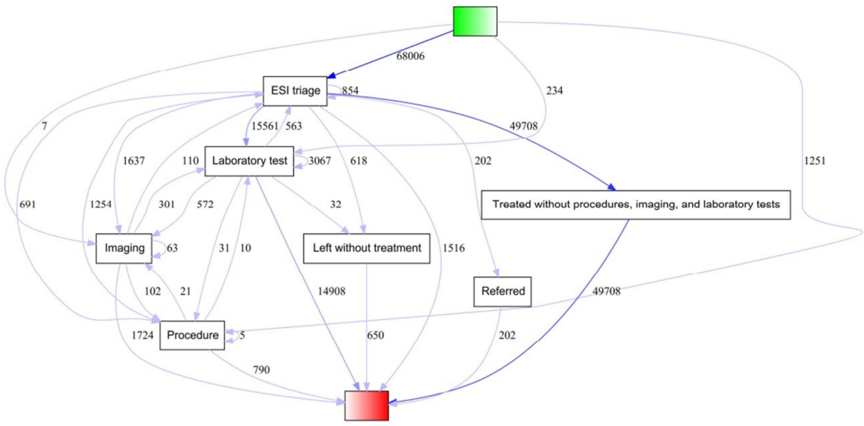

Figure 10: Process model for ESI 5 cluster.

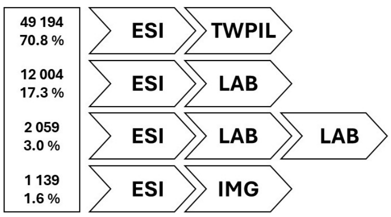

Figure 11: Top 4 pathways for ESI 5.

Visits	69 505
LOS (min)	83
Pro, Img, Lab events per visit	0.34
Trace variants	236
Visits of top 4 traces	64 396
Visits of top 4 traces (%)	92.6

Table 10: Key Metrics for ESI 5.

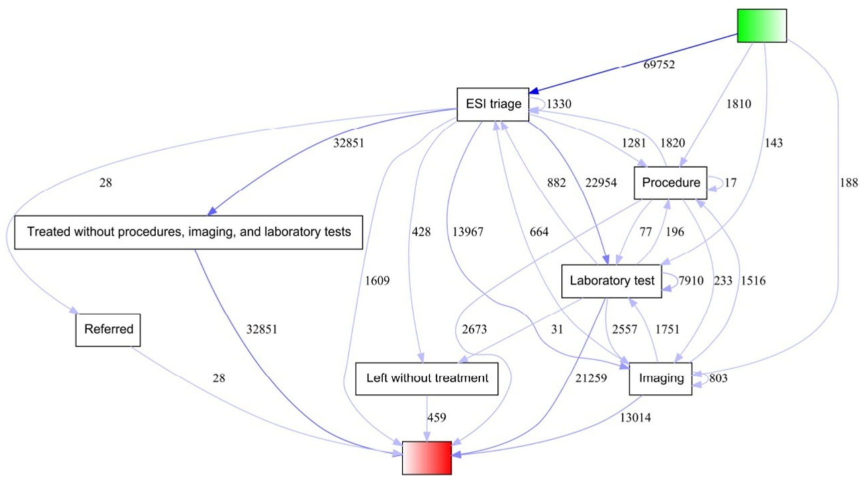

Figure 12: Process model for ESI 4 cluster.

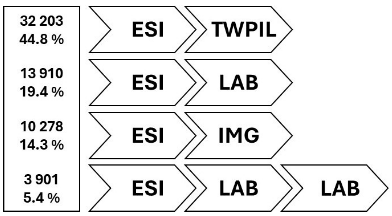

Figure 13: Top 4 pathways for ESI 4.

Visits	71 903
LOS (min)	135
Pro, Img, Lab events per visit	0.77
Trace variants	561
Visits of top 4 traces	60 292
Visits of top 4 traces (%)	83.9

Table 11: Key Metrics for ESI 4.

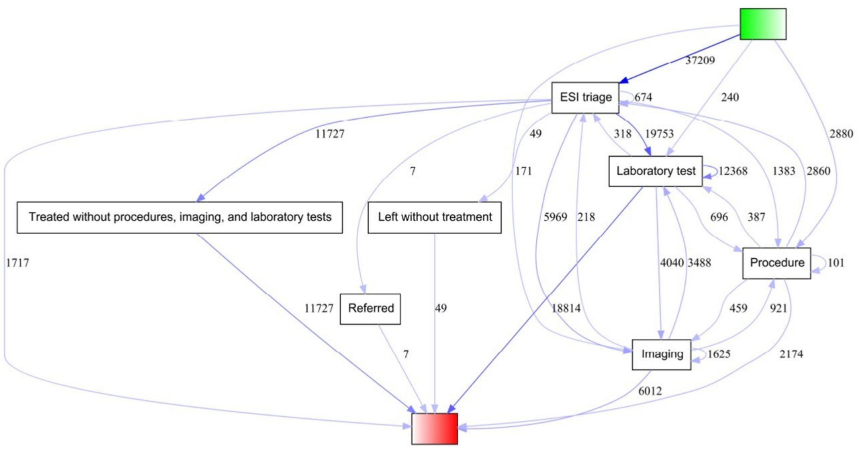

Figure 14: Process model for ESI 3 cluster.

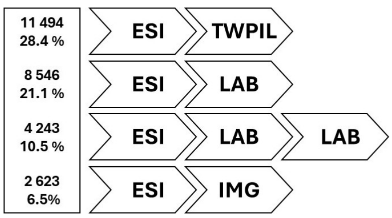

Figure 15: Top 4 pathways for ESI 3

Visits	40 506
LOS (min)	221
Pro, Img, Lab events per visit	1.3
Trace variants	669
Visits of top 4 traces	26 906
Visits of top 4 traces (%)	66.4

Table 12: Key Metrics for ESI 3.

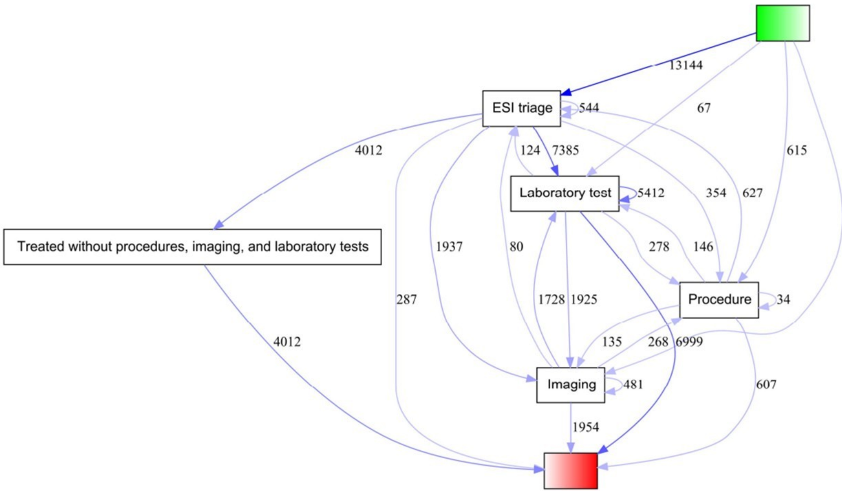

Figure 16: Process model for ESI 2 cluster.

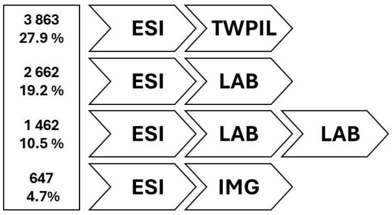

Figure 17: Top 4 pathways for ESI 2

Visits	13 867
LOS (min)	264
Pro, Img, Lab events per visit	1.5
Trace variants	456
Visits of top 4 traces	8 634
Visits of top 4 traces (%)	62.3

Table 13: Key Metrics for ESI 2.

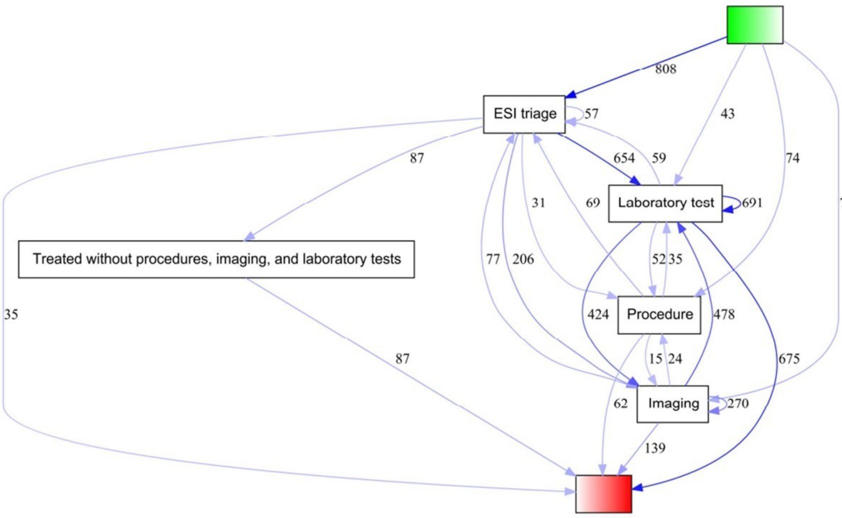

Figure 18: Process model for ESI 1 cluster.

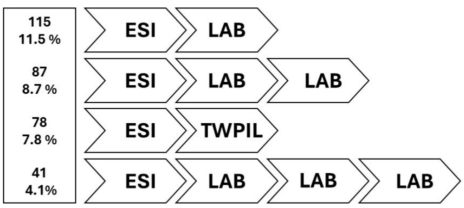

Figure 19: Top 4 pathways for ESI 1.

Visits	1 001
LOS (min)	270
Pro, Img, Lab events per visit	3.1
Trace variants	298
Visits of top 4 traces	321
Visits of top 4 traces (%)	32.1

Table 14: Key Metrics for ESI 1.

8 Discussion

8.1 Interpretation of Results

The general results obtained in this study align with consensus regarding pediatric and adolescent emergency departments. The patient cohort predominantly consists of non-urgent ED visits, where treatment typically does not involve procedures, imaging, or laboratory tests. As the focus shifts to more urgent ESI categories, the inclusion of procedures, imaging, and laboratory tests increases, leading to a greater demand for care and diagnostic resources. Nevertheless, the results obtained from the study should not be overlooked, as they provide valuable insights into the actual processes occurring in the emergency department. Furthermore, the results reveal a potential problem in the ESI triage assessment process and also in the ESI system itself.

A key finding from this study is that the four most common patient pathways account for 78 % of all emergency department visits at the pediatric and adolescent ED, suggesting that patients generally follow well-defined care pathways. However, since the four most frequent pathways require minimal resources in terms of procedures, imaging, and laboratory tests, this finding primarily reflects the fact that most patients either do not require these resources or need them only as a single instance. Consequently, no conclusions can be drawn about care processes that involve greater resource utilization.

The notion of well-defined patient pathways is contradicted by the process models in Figures 8, 10, 12, 14, 16, and 18, which reveal turbulent flow when the four most common traces are excluded. This is evident from the high-frequency bidirectional dependencies between the event classes for procedures, imaging, and laboratory tests, as well as the loops within themselves. Furthermore, the ratio of pathway variants relative to the total number of visits when excluding the four most frequent pathways points to the same conclusion. For ESI 5, this figure is 0.045, for ESI 4 it is 0.048, for ESI 3 it is 0.049, for ESI 2 it is 0.086, and for ESI 1 it is 0.43. In such a simplified model, where the primary drivers of variant formation are changes in the number and sequence of the three event classes (procedures, imaging, laboratory tests), these figure are relatively high.

This variation in patient pathways must be taken into consideration. While most patients are not affected by it, those who are tend to be more acute cases. It is essential to distinguish between variation arising from patient-related factors, such as acuity, differing medical needs, and background characteristics, and system-related factors, such as bottlenecks within the ED and differences in clinical decision-making among physicians and nursing staff. Identifying and understanding these sources of variability in future studies is crucial for improving efficiency and resource allocation in ED operations.

Before making conclusive judgments regarding the effectiveness of the ESI triage assessment process and the ESI system, it is important to acknowledge that the resource categorization employed in this study does not directly correspond to the resource categories defined by the ESI system. Simple and complex resources are not differentiated, consultations are not included in procedures as they are not structurally recorded in the HIS, and medication administration is not included. Additionally, POC laboratory tests, repeated laboratory tests, and imaging conducted with the same modality may have led to the creation of multiple resource events. Taking the potential error into account, the results still strongly indicate that the ESI triage assignment in the pediatric and adolescent EDs at NCH and Jorvi do not function as defined by the ESI system [69]. This is most evident in the ESI categories 5, 4, and 3, where the classification is based on resource utilization. A significant proportion of ESI 4 and ESI 3 patients are treated without procedures, imaging, or laboratory tests – approximately half and just under a third, respectively. Simultaneously, the resource events per visit are 0.77 and 1.3, respectively. Given that the ESI 4 definition specifies the need for one resource and the ESI 3 definition specifies the need for two or more, these figures demonstrate a clear misalignment

There are likely multiple factors contributing to this phenomenon. A significant contributor is the subjective and challenging nature of resource assessment. This is evident in the case of ESI 5 patients, where approximately 30 % ultimately undergo procedures, imaging, and laboratory tests, despite being initially assessed as not requiring them. In addition, physicians, particularly in low-acuity pediatric care, exhibit significant variability in their test ordering practices [89]. This variability cannot realistically or appropriately be accounted for during the ESI assignment phase due to the necessity of rapid and standardized triage decisions. Consequently, this may contribute to the observed misalignment between the estimated and actual resource requirements. Overall, it is not feasible to continuously assess resource needs in a manner that perfectly aligns with actual resource usage in the ED. Acceptance of some level of error is necessary.

Overcrowding, which is one of the major challenges in healthcare, is also likely one of the contributing factors to the ESI triage misalignment [2]. Overcrowding is defined as functional limitation caused by a patient influx that surpasses the available space and staffing capacity [90]. It is a significant issue across the healthcare system, but it is particularly evident in pediatric care, where it manifests as patients coming to the ED to address non-urgent issues [91]. In addition to contributing to ED overcrowding, this may also impact the alignment of the ESI triage system. If the emergency department is filled with non-urgent issues, ideally not requiring emergency care but categorized as ESI 5, the threshold for assigning ESI 4 might be lowered. This to ensure that a patient requiring treatment within a hours, who would normally be classified as ESI 5, receives care before truly non-urgent cases. On the other hand, this creates additional pressure to assign ESI 4 level patients to the ESI 3 category. In a sense, this means that the ESI scale is adjusted downward to accommodate non-urgent issues within the triage system.

It is also important to acknowledge that the ESI triage system is not without its limitations. The relationship between the required ESI resources and the clinical urgency of a condition is not always consistent. A notable example is psychiatric emergencies, which may demand rapid or immediate intervention without requiring any ESI resources [92]. Likewise, conditions such as asthma exacerbations and hypoglycemia, while necessitating urgent care similar to ESI 3 level urgency, might require no more than one ESI resource [93]. This is further supported by the comparable numbers of patient pathways in the ESI 3 and ESI 2 categories that do not involve procedures, imaging, and laboratory tests. Since the ESI 2 level is based on the acuity of patient, and the ratio of patient pathways for ESI 3 that do not involve procedures, imaging, and laboratory tests is the same as for ESI 2, it is unlikely that all such ESI 2 cases are acute, while all ESI 3 cases are incorrectly triaged. On the other hand, based on this study, it cannot be definitively concluded how many of these ESI 2 cases have been correctly triaged.

This study does not allow for determining the extent to which different factors contribute to the misalignment between the ESI triage system and the actual pathways observed in each ESI triage category, making it challenging to assess the consequences at an individual level. However, regardless of the contributing factors to the misalignment, this remains a significant issue at the systemic level. For operational management, the functionality of triage is critically important. While triage was primarily developed to assess the urgency of patient care needs, its other key role is to ensure that departmental resources are utilized most efficiently [94, 68]. When triage functions differently than intended, achieving this latter goal becomes particularly challenging.

There are two main reasons for this: immediate resource allocation and future resource planning. In either case, whether the observed misalignment is due to patients being treated in too low ESI categories (more urgent) relative to their condition, or patients being treated with the correct urgency but misclassified according to the ESI triage system due to its limitations, problems arise in both scenarios. In the first scenario, there are likely direct effects on the overuse of resources relative to the condition. These can include, for example, treatment in specialized care rather than primary care, occupation of inpatient beds, and closer monitoring. The immediate effects of the second scenario are much harder to assess, as the lack of a standardized system means there is no benchmark for comparison.

In a broader context, both scenarios present significant challenges for overall operational management due to issues with data and performance measurement. This is because evidence-based management heavily relies on data. If there are indications that the data reflects an inefficient process, but the extent of the inefficiency or the causes behind it cannot be determined, the value of the data as a performance metric diminishes. In this case, improving the process efficiently becomes more of a trial-and-error approach rather than being evidence-based. Therefore, it is crucial, for patient safety but also for process efficiency, that triage functions according to established standards. This is an area that must be investigated further to determine the true level of triage misalignment and the underlying causes.

Lastly, the process models show that nearly half of the procedures are recorded as being performed before the ESI triage assessment. This is most clearly seen in the overall emergency department process presented in Figure 8. This raises the question of whether these recorded procedures were actually performed before or at the same time as the ESI triage assessment, or if they are a result of time stamping errors in the HIS system or a systematic issue with the staff. For studies like this and the development of evidence-based medicine, it is crucial that all HIS timestamps are accurate, and this requires further investigation.

8.2 Limitations of Study

This study has several limitations. Firstly, due to its retrospective nature, there is a potential for errors in data sourcing or documentation. The trade-offs involved in data filtering and event log structuring, as discussed in sections 6.1.1 and 6.1.2, may also have a potential impact on the results. This is typical of retrospective statistical studies. Secondly, because ESI clusters differ significantly in visit numbers, particularly for ESI 1, smaller clusters may show relatively higher trace variability compared to larger clusters. On the other hand, comparing the relative proportion of the four most common patient pathways across all pathways and comparing these across ESI categories helps reduce bias compared to directly comparing trace variant ratios. Thirdly, as stated in section 8.1, the resources for procedures, imaging, and laboratory tests in this study are not directly comparable to ESI resources. The purpose of this study was not to develop ESI resource models but rather to process mine the pathways of pediatric and adolescent ED patients. The process models created can provide insights into the functioning of the ESI triage system.

In the broader context, this study does not comprehensively capture the entire patient pathway within the emergency department. For instance, it omits events like assessments and examinations conducted by doctors and nurses. Additionally, medication annotations are excluded from the study, potentially leading to gaps in the overall documentation of the patient pathways. Additionally, the study reflects the typical limitations associated with DFM Miner, such as the process models generated being constrained in that they only capture immediate dependencies between actions. However, despite its limitations, this study represents a pioneering effort in the application of process mining at this scale, both within the context of pediatric and adolescent emergency care as well as general emergency departments. While there are constraints, its significance lies in laying the groundwork for future study in this area.

8.3 Applicability of Process Mining in Study

This study has shown that process mining can be effectively applied to complex healthcare settings with large datasets. The DFG miner demonstrates its utility in identifying pathway variability, resource allocation, and potential failure demand. The data driven process model clearly shows the pathways patients follow in the emergency department and is more intuitive to comprehend than just using process metrics. Its straightforward algorithm succeeds in creating a clear and comprehensible process model from complex and extensive healthcare data, where other miners fail. However, several significant limitations remain with the DFG Miner. These include its restricted ability to capture only immediate dependencies between actions, as already mentioned in Section 8.2, and the broader limitations described in Section 7. These limitations mean that while patient pathway variability, resource allocation, and potential failure demand can be identified, the underlying causes can only be speculated.

Since identifying the consequences alone does not provide practical benefits, it is crucial to move to the next stage by identifying the underlying causes through process mining. There are three approaches to leveraging process mining for this purpose. The first approach is to break the data down into smaller clusters, such as by outcome or specialty, and use DFG or potentially other existing miners to identify the specific characteristics of these clusters. This approach can help identify specific characteristics between clusters and uncover the underlying causes for the observed pathway variability, resource allocation, and potential failure demand. This has already been initiated as part of the continuation of this project. However, the method is labor-intensive, and while it provides valuable insights into different patient groups, it remains uncertain how effectively this information contributes to the overall view.

The second approach is to develop current miners with custom algorithms. There is a significant demand for an effective and scalable process miner that works efficiently with healthcare data. Assessing the feasibility of implementing this is challenging. The third approach is to combine process mining with machine learning and other data science disciplines to create an entirely new method for process mining. This approach has the potential to revolutionize healthcare process mining with its unparalleled ability to link variables together and distinguish cause-and-effect relationships. However, to date, there is no widely adopted tool that effectively combines these techniques.

9 Conclusions

Process mining is a powerful method for enhancing data-driven understanding of real-world processes. Its rapid development in the 21st century has greatly expanded its applications. It is now widely used in the academic world and is increasingly being applied in the business sector. However, several limitations still hinder its broader adoption in operational contexts, particularly within healthcare.

This study presents novel research on the application of current process mining tools to large scale healthcare datasets. The findings demonstrate that the DFM miner is effective in analyzing these datasets and generating valuable insights. Despite the simplicity of the DFM miner, it should not be underestimated as an effective tool for creating accurate models from complex event logs. The results further indicate that, in pediatric emergency departments, the majority of patients require minimal resource utilization. Additionally, the effectiveness of the ESI triage system remains inconclusive. At the same time, the study emphasizes the need for continued advancements in process mining methodologies to fully exploit their potential in healthcare settings. While current methods can reveal what occurs within the emergency department, they offer limited insight into the underlying causes of these processes.

This study has already generated significant interest within HUS and is expected to serve as a benchmark for future research. Based on the findings, follow-up studies and algorithm development are already underway. It is crucial to develop methods to assess the true impact of various factors on emergency department operations and evaluate the effectiveness of existing practices, such as the ESI triage system. This requires expanding the scope to include factors such as human resources, emergency department occupancy rates, patient vitals, and natural variables like time of day. To address such a complex challenge, the scope of process mining should extend beyond its current boundaries to incorporate other data science disciplines.

These efforts aim to enhance the impact and efficiency of healthcare. This study represents the first step toward the broader adoption of process mining at HUS. The ultimate goal is to establish a solid foundation for the role of process mining in evidence-based medicine and management, contributing to improved healthcare outcomes and greater operational efficiency.

References

References

[1] R. S. Mans, W. v. d. Aalst, and R. J. Vanwersch, *Process Mining in Healthcare: Evaluating and Exploiting Operational Healthcare Processes*. Springer, 2015.

[2] G. Savioli, I. F. Ceresa, N. Gri, G. Bavestrello Piccini, Y. Longhitano, C. Zanza, A. Piccioni, C. Esposito, G. Ricevuti, and M. A. Bressan, "Emergency department overcrowding: understanding the factors to find corresponding solutions," *Journal of personalized medicine*, vol. 12, no. 2, p. 279, 2022.

[3] A. D'Andreamatteo, L. Ianni, F. Lega, and M. Sargiacomo, "Lean in healthcare: A comprehensive review," *Health policy*, vol. 119, no. 9, pp. 1197–1209, 2015.

[4] M. Netjes, R. S. Mans, H. A. Reijers, W. M. van der Aalst, and R. J. Vanwersch, "Bpr best practices for the healthcare domain," in *Business Process Management Workshops: BPM 2009 International Workshops, Ulm, Germany, September 7, 2009. Revised Papers 7*, pp. 605–616, Springer, 2010.

[5] K. Walshe and T. G. Rundall, "Evidence-based management: from theory to practice in health care," *The Milbank Quarterly*, vol. 79, no. 3, pp. 429–457, 2001.

[6] J. Munoz-Gama, N. Martin, C. Fernandez-Llatas, O. A. Johnson, M. Sepúlveda, E. Helm, V. Galvez-Yanjari, E. Rojas, A. Martinez-Millana, D. Aloini, *et al.*, "Process mining for healthcare: Characteristics and challenges," *Journal of Biomedical Informatics*, vol. 127, p. 103994, 2022.

[7] E. Teisberg, S. Wallace, and S. O'Hara, "Defining and implementing value-based health care: a strategic framework," *Academic Medicine*, vol. 95, no. 5, pp. 682–685, 2020.

[8] M. Christ, F. Grossmann, D. Winter, R. Bingisser, and E. Platz, "Modern triage in the emergency department," *Deutsches Ärzteblatt International*, vol. 107, no. 50, p. 892, 2010.

[9] P. Parviainen, M. Tihinen, J. Kääriäinen, and S. Teppola, "Tackling the digitalization challenge: how to benefit from digitalization in practice," *International journal of information systems and project management*, vol. 5, no. 1, pp. 63–77, 2017.

[10] D. Blazquez and J. Domenech, "Big data sources and methods for social and economic analyses," *Technological Forecasting and Social Change*, vol. 130, pp. 99–113, 2018.

[11] W. M. Van der Aalst, “Data scientist: The engineer of the future,” in *Enterprise interoperability VI: Interoperability for agility, resilience and plasticity of collaborations*, pp. 13–26, Springer, 2014.

[12] T. Ramakrishnan, M. C. Jones, and A. Sidorova, “Factors influencing business intelligence (bi) data collection strategies: An empirical investigation,” *Decision support systems*, vol. 52, no. 2, pp. 486–496, 2012.

[13] B. List, R. M. Bruckner, K. Machaczek, and J. Schiefer, “A comparison of data warehouse development methodologies case study of the process warehouse,” in *Database and Expert Systems Applications: 13th International Conference, DEXA 2002 Aix-en-Provence, France, September 2–6, 2002 Proceedings 13*, pp. 203–215, Springer, 2002.

[14] S. Negash and P. Gray, “Business intelligence,” *Handbook on decision support systems 2*, pp. 175–193, 2008.

[15] W. Van Der Aalst and W. van der Aalst, *Data science in action*. Springer, 2016.

[16] W. M. Van der Aalst, J. Nakatumba, A. Rozinat, and N. Russell, “Business process simulation,” *Handbook on Business Process Management 1: Introduction, Methods, and Information Systems*, pp. 313–338, 2010.

[17] S. Suriadi, R. Andrews, A. H. ter Hofstede, and M. T. Wynn, “Event log imperfection patterns for process mining: Towards a systematic approach to cleaning event logs,” *Information systems*, vol. 64, pp. 132–150, 2017.

[18] W. M. Van der Aalst, “Extracting event data from databases to unleash process mining,” in *BPM-Driving innovation in a digital world*, pp. 105–128, Springer, 2015.

[19] H. Verbeek, J. C. Buijs, B. F. Van Dongen, and W. M. Van Der Aalst, “Xes, xesame, and prom 6,” in *Information Systems Evolution: CAiSE Forum 2010, Hammamet, Tunisia, June 7-9, 2010, Selected Extended Papers 22*, pp. 60–75, Springer, 2011.

[20] W. M. van der Aalst, “Process discovery: Capturing the invisible,” *IEEE Computational Intelligence Magazine*, vol. 5, no. 1, pp. 28–41, 2010.

[21] W. van der Aalst, A. Adriansyah, A. K. A. de Medeiros, F. Arcieri, T. Baier, T. Blickle, J. C. Bose, P. van den Brand, R. Brandtjen, J. Buijs, A. Burattin, J. Carmona, M. Castellanos, J. Claes, J. Cook, N. Costantini, F. Curbera, E. Damiani, M. de Leoni, P. Delias, B. F. van Dongen, M. Dumas, S. Dustdar, D. Fahland, D. R. Ferreira, W. Gaaloul, F. van Geffen, S. Goel, C. Günther, A. Guzzo, P. Harmon, A. ter Hofstede, J. Hoogland, J. E. Ingvaldsen, K. Kato, R. Kuhn, A. Kumar, M. La Rosa, F. Maggi, D. Malerba, R. S. Mans, A. Manuel, M. McCreesh, P. Mello, J. Mendling, M. Montali, H. R. Motahari-Nezhad, M. zur Muehlen, J. Munoz-Gama, L. Pontieri, J. Ribeiro,

A. Rozinat, H. Seguel Pérez, R. Seguel Pérez, M. Sepúlveda, J. Sinur, P. Soffer, M. Song, A. Sperduti, G. Stilo, C. Stoel, K. Swenson, M. Talamo, W. Tan, C. Turner, J. Vanthienen, G. Varvaressos, E. Verbeek, M. Verdonk, R. Vigo, J. Wang, B. Weber, M. Weidlich, T. Weijters, L. Wen, M. Westergaard, and M. Wynn, “"process mining manifesto",” in *"Business Process Management Workshops"*, pp. "169–194", "Springer Berlin Heidelberg", "2012".

[22] W. Reisig, *Petri nets: an introduction*, vol. 4. Springer Science & Business Media, 2012.

[23] M. Chinosi and A. Trombetta, “Bpmn: An introduction to the standard,” *Computer Standards & Interfaces*, vol. 34, no. 1, pp. 124–134, 2012.

[24] W. M. van der Aalst, “A practitioner's guide to process mining: Limitations of the directly-follows graph,” *Procedia Computer Science*, vol. 164, pp. 321–328, 2019. CENTERIS 2019 - International Conference on ENTERprise Information Systems / ProjMAN 2019 - International Conference on Project MANagement / HCist 2019 - International Conference on Health and Social Care Information Systems and Technologies, CENTERIS/ProjMAN/HCist 2019.

[25] A. Rozinat and W. M. Van der Aalst, “Conformance checking of processes based on monitoring real behavior,” *Information Systems*, vol. 33, no. 1, pp. 64–95, 2008.

[26] M. De Leoni, F. M. Maggi, and W. M. van der Aalst, “Aligning event logs and declarative process models for conformance checking,” in *Business Process Management: 10th International Conference, BPM 2012, Tallinn, Estonia, September 3-6, 2012. Proceedings 10*, pp. 82–97, Springer, 2012.

[27] W. Van Der Aalst, “Process mining: Overview and opportunities,” *ACM Transactions on Management Information Systems (TMIS)*, vol. 3, no. 2, pp. 1–17, 2012.

[28] S. Mukherjee and A. Chatterjee, “The concept of bottleneck,” in *International Conference on Multi-Echelon*, pp. 37–48, 2007.

[29] A. Rozinat and W. M. van der Aalst, “Decision mining in prom,” in *Business Process Management: 4th International Conference, BPM 2006, Vienna, Austria, September 5-7, 2006. Proceedings 4*, pp. 420–425, Springer, 2006.

[30] E. Bazhenova and M. Weske, “Deriving decision models from process models by enhanced decision mining,” in *Business Process Management Workshops: BPM 2015, 13th International Workshops, Innsbruck, Austria, August 31–September 3, 2015, Revised Papers 14*, pp. 444–457, Springer, 2016.

[31] W. Van Der Aalst, “Process mining,” *Communications of the ACM*, vol. 55, no. 8, pp. 76–83, 2012.

[32] J. S. Cristóbal, "Complexity in project management," *Procedia Computer Science*, vol. 121, pp. 762–766, 2017. CENTERIS 2017 - International Conference on ENTERprise Information Systems / ProjMAN 2017 - International Conference on Project MANagement / HCist 2017 - International Conference on Health and Social Care Information Systems and Technologies, CENTERIS/ProjMAN/HCist 2017.

[33] S. LaValle, E. Lesser, R. Shockley, M. S. Hopkins, and N. Kruschwitz, "Big data, analytics and the path from insights to value," *MIT sloan management review*, 2010.

[34] H.-J. Cheng and A. Kumar, "Process mining on noisy logs—can log sanitization help to improve performance?," *Decision Support Systems*, vol. 79, pp. 138–149, 2015.

[35] S. J. Leemans, E. Poppe, and M. T. Wynn, "Directly follows-based process mining: Exploration & a case study," in *2019 International Conference on Process Mining (ICPM)*, pp. 25–32, IEEE, 2019.

[36] W. Van der Aalst, T. Weijters, and L. Maruster, "Workflow mining: Discovering process models from event logs," *IEEE transactions on knowledge and data engineering*, vol. 16, no. 9, pp. 1128–1142, 2004.

[37] A. J. Weijters, W. M. van Der Aalst, and A. A. De Medeiros, "Process mining with the heuristicsminer algorithm," 2006.

[38] C. W. Günther and W. M. Van Der Aalst, "Fuzzy mining–adaptive process simplification based on multi-perspective metrics," in *International conference on business process management*, pp. 328–343, Springer, 2007.

[39] A. Adriansyah, B. F. van Dongen, and W. M. van der Aalst, "Conformance checking using cost-based fitness analysis," in *2011 ieee 15th international enterprise distributed object computing conference*, pp. 55–64, IEEE, 2011.

[40] W. M. Van Der Aalst, H. A. Reijers, and M. Song, "Discovering social networks from event logs," *Computer Supported Cooperative Work (CSCW)*, vol. 14, pp. 549–593, 2005.

[41] M. Song and W. M. Van der Aalst, "Towards comprehensive support for organizational mining," *Decision support systems*, vol. 46, no. 1, pp. 300–317, 2008.

[42] F. Mannhardt, M. De Leoni, and H. A. Reijers, "Heuristic mining revamped: an interactive, data-aware, and conformance-aware miner," in *15th International Conference on Business Process Management (BPM 2017)*, pp. 1–5, CEUR-WS. org, 2017.

[43] J. Pang, H. Xu, J. Ren, J. Yang, M. Li, D. Lu, and D. Zhao, "Process mining framework with time perspective for understanding acute care: a case study of ais in hospitals," *BMC Medical Informatics and Decision Making*, vol. 21, pp. 1–10, 2021.

[44] J. E. Cook and A. L. Wolf, "Event-based detection of concurrency," *ACM SIGSOFT Software Engineering Notes*, vol. 23, no. 6, pp. 35–45, 1998.

[45] W. M. van der Aalst, R. De Masellis, C. Di Francescomarino, and C. Ghidini, "Learning hybrid process models from events: Process discovery without faking confidence," in *Business Process Management: 15th International Conference, BPM 2017, Barcelona, Spain, September 10–15, 2017, Proceedings 15*, pp. 59–76, Springer, 2017.

[46] W. M. Van Der Aalst, "A practitioner's guide to process mining: Limitations of the directly-follows graph," 2019.

[47] A. Berti, S. J. Van Zelst, and W. van der Aalst, "Process mining for python (pm4py): bridging the gap between process-and data science," *arXiv preprint arXiv:1905.06169*, 2019.

[48] B. F. Van Dongen, A. K. A. de Medeiros, H. Verbeek, A. Weijters, and W. M. van Der Aalst, "The prom framework: A new era in process mining tool support," in *Applications and Theory of Petri Nets 2005: 26th International Conference, ICATPN 2005, Miami, USA, June 20-25, 2005. Proceedings 26*, pp. 444–454, Springer, 2005.

[49] A. Berti, S. van Zelst, and D. Schuster, "Pm4py: a process mining library for python," *Software Impacts*, vol. 17, p. 100556, 2023.

[50] A. Berti, S. van Zelst, and D. Schuster, "Pm4py: A process mining library for python," *Software Impacts*, vol. 17, p. 100556, 2023.

[51] T. E. Oliphant *et al.*, *Guide to numpy*, vol. 1. Trelgol Publishing USA, 2006.

[52] W. McKinney *et al.*, "pandas: a foundational python library for data analysis and statistics," *Python for high performance and scientific computing*, vol. 14, no. 9, pp. 1–9, 2011.

[53] A. Gulli and S. Pal, *Deep learning with Keras*. Packt Publishing Ltd, 2017.

[54] C. W. Günther and A. Rozinat, "Disco: Discover your processes," in *Demonstration Track of the 10th International Conference on Business Process Management, BPM Demos 2012*, pp. 40–44, CEUR-WS. org, 2012.

[55] C. J. Turner, A. Tiwari, R. Olaiya, and Y. Xu, "Process mining: from theory to practice," *Business process management journal*, vol. 18, no. 3, pp. 493–512, 2012.

[56] U. Celik and E. Akçetin, “Process mining tools comparison,” *Online Academic Journal of Information Technology*, vol. 9, no. 34, pp. 97–104, 2018.

[57] P. Drakoulogkonas and D. Apostolou, “On the selection of process mining tools,” *Electronics*, vol. 10, no. 4, p. 451, 2021.

[58] A. F. D. Gomes, C. Wanzeller, and J. Fialho, “Comparative analysis of process mining tools,” 2021.

[59] C. Parente and C. J. Costa, “Comparing process mining tools and algorithms,” in *2022 17th Iberian Conference on Information Systems and Technologies (CISTI)*, pp. 1–7, IEEE, 2022.

[60] W. J. Trybula, “Data mining and knowledge discovery.,” *Annual review of information science and technology (ARIST)*, vol. 32, pp. 197–229, 1997.

[61] M.-S. Chen, J. Han, and P. S. Yu, “Data mining: an overview from a database perspective,” *IEEE Transactions on Knowledge and data Engineering*, vol. 8, no. 6, pp. 866–883, 1996.

[62] E. De Roock and N. Martin, “Process mining in healthcare–an updated perspective on the state of the art,” *Journal of biomedical informatics*, vol. 127, p. 103995, 2022.

[63] M. R. Dallagassa, C. dos Santos Garcia, E. E. Scalabrin, S. O. Ioshii, and D. R. Carvalho, “Opportunities and challenges for applying process mining in healthcare: a systematic mapping study,” *Journal of Ambient Intelligence and Humanized Computing*, pp. 1–18, 2022.

[64] A. Harvey, A. Brand, S. T. Holgate, L. V. Kristiansen, H. Lehrach, A. Palotie, and B. Prainsack, “The future of technologies for personalised medicine,” *New biotechnology*, vol. 29, no. 6, pp. 625–633, 2012.

[65] N. Martin and J. Bergs, “Patient flow data registration: A key barrier to the data-driven and proactive management of an emergency department,” *International Emergency Nursing, 53*, 2020.

[66] S. Alter, “Theory of workarounds,” 2014.

[67] Tilastokeskus, “Statfin: Väestö- ja elinolo -statistiikka.” https://pxdata.stat.fi/PxWeb/pxweb/fi/StatFin/StatFin__vaerak/, 2024. Accessed: 2024-10-28.

[68] K. V. Iserson and J. C. Moskop, “Triage in medicine, part i: concept, history, and types,” *Annals of emergency medicine*, vol. 49, no. 3, pp. 275–281, 2007.

[69] N. Gilboy, P. Tanabe, D. Travers, and A. M. Rosenau, *Emergency Severity Index (ESI): A Triage Tool for Emergency Department Care, Version 4. Implementation Handbook 2012 Edition*. Rockville, MD: Agency for Healthcare Research and Quality, November 2011. AHRQ Publication No. 12-0014.

[70] R. Shelton, "The emergency severity index 5-level triage system," *Dimensions of Critical Care Nursing*, vol. 28, no. 1, pp. 9–12, 2009.

[71] E. Rojas, J. Munoz-Gama, M. Sepúlveda, and D. Capurro, "Process mining in healthcare: A literature review," *Journal of biomedical informatics*, vol. 61, pp. 224–236, 2016.

[72] M. Ghasemi and D. Amyot, "Process mining in healthcare: a systematised literature review," *International Journal of Electronic Healthcare*, vol. 9, no. 1, pp. 60–88, 2016.

[73] T. G. Erdogan and A. Tarhan, "Systematic mapping of process mining studies in healthcare," *IEEE Access*, vol. 6, pp. 24543–24567, 2018.

[74] C. dos Santos Garcia, A. Meincheim, E. R. F. Junior, M. R. Dallagassa, D. M. V. Sato, D. R. Carvalho, E. A. P. Santos, and E. E. Scalabrin, "Process mining techniques and applications–a systematic mapping study," *Expert Systems with Applications*, vol. 133, pp. 260–295, 2019.

[75] A. B. Durojaiye, N. M. McGeorge, L. L. Puett, D. Stewart, J. C. Fackler, P. L. Hoonakker, H. P. Lehmann, and A. P. Gurses, "Mapping the flow of pediatric trauma patients using process mining," *Applied clinical informatics*, vol. 9, no. 03, pp. 654–666, 2018.

[76] A. B. Durojaiye, S. Levin, M. Toerper, H. Kharrazi, H. P. Lehmann, and A. P. Gurses, "Evaluation of multidisciplinary collaboration in pediatric trauma care using ehr data," *Journal of the American Medical Informatics Association*, vol. 26, no. 6, pp. 506–515, 2019.

[77] R. C. Basole, M. L. Braunstein, V. Kumar, H. Park, M. Kahng, D. H. Chau, A. Tamersoy, D. A. Hirsh, N. Serban, J. Bost, *et al.*, "Understanding variations in pediatric asthma care processes in the emergency department using visual analytics," *Journal of the American Medical Informatics Association*, vol. 22, no. 2, pp. 318–323, 2015.

[78] D. Duma and R. Aringhieri, "An ad hoc process mining approach to discover patient paths of an emergency department," *Flexible Services and Manufacturing Journal*, vol. 32, pp. 6–34, 2020.

[79] F. Rismanchian and Y. H. Lee, "Process mining–based method of designing and optimizing the layouts of emergency departments in hospitals," *HERD: Health Environments Research & Design Journal*, vol. 10, no. 4, pp. 105–120, 2017.

[80] R. Andrews, S. Suriadi, M. Wynn, A. H. ter Hofstede, and S. Rothwell, "Improving patient flows at st. andrew's war memorial hospital's emergency department through process mining," *Business Process Management Cases: Digital Innovation and Business Transformation in Practice*, pp. 311–333, 2018.

[81] C. Alvarez, E. Rojas, M. Arias, J. Munoz-Gama, M. Sepúlveda, V. Herskovic, and D. Capurro, "Discovering role interaction models in the emergency room using process mining," *Journal of biomedical informatics*, vol. 78, pp. 60–77, 2018.

[82] M. Cho, M. Song, J. Park, S.-R. Yeom, I.-J. Wang, and B.-K. Choi, "Process mining-supported emergency room process performance indicators," *International journal of environmental research and public health*, vol. 17, no. 17, p. 6290, 2020.

[83] T. G. Erdogan and A. K. Tarhan, "Multi-perspective process mining for emergency process," *Health informatics journal*, vol. 28, no. 1, p. 14604582221077195, 2022.

[84] E. Rojas, A. Cifuentes, A. Burattin, J. Munoz-Gama, M. Sepúlveda, and D. Capurro, "Performance analysis of emergency room episodes through process mining," *International journal of environmental research and public health*, vol. 16, no. 7, p. 1274, 2019.

[85] F. Davari, M. Nasr Isfahani, A. Atighechian, and E. Ghobadian, "Optimizing emergency department efficiency: a comparative analysis of process mining and simulation models to mitigate overcrowding and waiting times," *BMC Medical Informatics and Decision Making*, vol. 24, no. 1, p. 295, 2024.

[86] G. Ibanez-Sanchez, C. Fernandez-Llatas, A. Martinez-Millana, A. Celda, J. Mandingorra, L. Aparici-Tortajada, Z. Valero-Ramon, J. Munoz-Gama, M. Sepúlveda, E. Rojas, *et al.*, "Toward value-based healthcare through interactive process mining in emergency rooms: the stroke case," *International journal of environmental research and public health*, vol. 16, no. 10, p. 1783, 2019.

[87] N. Martin, "Data quality in process mining," *Interactive process mining in healthcare*, pp. 53–79, 2021.

[88] A. Pika, M. T. Wynn, S. Budiono, A. H. Ter Hofstede, W. M. van der Aalst, and H. A. Reijers, "Privacy-preserving process mining in healthcare," *International journal of environmental research and public health*, vol. 17, no. 5, p. 1612, 2020.

[89] K. Doctor, K. Breslin, J. M. Chamberlain, and D. Berkowitz, "Practice pattern variation in test ordering for low-acuity pediatric emergency department patients," *Pediatric emergency care*, vol. 37, no. 3, pp. e116–e123, 2021.

[90] A. C. for Emergency Medicine, "Policy on standard terminology," 2023. Accessed: 2024-11-28.

[91] M. Sartini, A. Carbone, A. Demartini, L. Giribone, M. Oliva, A. M. Spagnolo, P. Cremonesi, F. Canale, and M. L. Cristina, "Overcrowding in emergency department: causes, consequences, and solutions—a narrative review," in *Healthcare*, vol. 10, p. 1625, MDPI, 2022.

[92] L. Zun, "Care of psychiatric patients: the challenge to emergency physicians," *Western journal of emergency medicine*, vol. 17, no. 2, p. 173, 2016.

[93] H. Jang, M. Ozkaynak, T. Ayer, and M. R. Sills, "Factors associated with first medication time for children treated in the emergency department for asthma," *Pediatric Emergency Care*, vol. 37, no. 1, pp. e42–e47, 2021.

[94] G. FitzGerald, G. A. Jelinek, D. Scott, and M. F. Gerdtz, "Emergency department triage revisited," *Emergency Medicine Journal*, vol. 27, no. 2, pp. 86–92, 2010.