Process Mining Using Electronic Health Records Data – Quo Vadis? Reflections from Observing Nurses' Activities and Data Registration Behavior

Niels Martin
UHasselt - Hasselt University
niels.martin@uhasselt.be

Isabeau Gielen
UHasselt - Hasselt University
isabeau.gielen@student.uhasselt.be

Jochen Bergs
UHasselt - Hasselt University
jochen.bergs@uhasselt.be

Abstract

Process mining leverages process execution data to better understand and improve operational processes. In hospitals, data from the Electronic Health Records (EHR) system that supports their daily operations is often used as input data for process mining. As limitations of EHR data in terms of data quality have also been highlighted in literature, it remains an open question how well EHR data reflects how work actually gets done in a care process. Against this background, this paper reports on the outcomes of an observation study at a Belgian hospital. In particular, the activities that nurses perform have been observed, as well as their data registration behavior. From the findings, it follows that EHR data will provide a highly fragmented and inaccurate view of how nursing work gets done. This constitutes a basis for reflection upon the extent to which EHR data is a truthful basis for process mining.

Keywords: process mining, healthcare, EHR data, observations, data registration behavior

1. Introduction

Process mining provides a rich collection of methods that leverage an event log to better understand and improve operational processes (Munoz-Gama et al., 2022; van der Aalst, 2022). An event log, the key input for process mining, consists of data reflecting the real-life execution of a process, which can substantially deviate from how process participants perceive the process, demonstrating the power of process mining (van der Aalst, 2016). Over the past two decades, process mining methods have been developed for a wide variety of use cases such as deriving a control-flow model that visualizes the order of activities, assessing the extent to which the real-life execution of a process deviates from a normative model, pinpointing the bottlenecks in a process, and predicting what the next activity for a running process instance will be (van der Aalst, 2016; van der Aalst, 2022). Healthcare is one of the prominent application areas in the process mining research field (De Roock & Martin, 2022; Zerbino et al., 2021). While the healthcare sector encompasses a wide range of healthcare organizations (*e.g.,* home healthcare organizations and elderly care organizations) the majority of works on process mining in healthcare target hospitals, which will also be the prime focus of this paper.

An event log is constructed from process execution data that is captured by an information system (van der Aalst, 2016). Within a hospital setting, an increasing adoption of Electronic Health Record (EHR) systems is observed (Adler-Milstein et al., 2015; Poba-Nzaou & Uwizeyemungu, 2019), which are systems that support planning of patient care, documenting how this care is delivered, and assessing the care outcomes (Häyrinen et al., 2008). While these systems are designed to support the daily hospital's operations (*e.g.,* by making the administered medication visible), the data it records is often used as a starting point for process mining as it contains snippets of data that reflect how care processes are being performed (Kusuma et al., 2021). For instance: the EHR will automatically keep track of the time at which an entry is made in a patient's file, which is a key ingredient of an event log (Rule et al., 2020; van der Aalst, 2016).

Despite the potential of EHR data as a source of input data for process mining, existing literature has also acknowledged its limitations. In particular, a wide variety of data quality issues have been reported (Fox et al., 2018; Martin, 2021; Munoz-Gama et al.,

URI: https://hdl.handle.net/10125/106830
978-0-9981331-7-1

2022; Vanbrabant et al., 2019). A prime example of such a data quality issue is a mismatch between the time and potentially even the order at which care activities are executed on the one hand, and the moment and order at which it is recorded in the EHR on the other hand (Martin, 2021; Martin et al., 2022). Other examples include missing events (*i.e.,* activities that are not recorded at all while they were performed in reality), imprecise timestamps (*e.g.,* only dates are available instead of more fine-grained timestamps), and imprecise resource data (*e.g.,* the data only reflects a department associated with a specific activity, but not the specific resource) (Mans et al., 2015).

As process mining centers around the premise that it provides in-depth insights into the real-life execution of a process from data (van der Aalst, 2022), the reported data quality issues raise the question of how well EHR data reflects how work actually gets done in a care process. While the presence of these data quality issues is widely acknowledged by the research community (Munoz-Gama et al., 2022), surprisingly little attention is attributed to better understanding the extent to which EHR data is a truthful starting point for process mining altogether. To contribute to such an understanding, this paper reports on the outcomes of an observation study at a residential ward of a Belgian hospital. In particular, the activities that nurses perform have been observed, as well as their data registration behavior (*i.e.,* when is data inserted into the EHR and which data is recorded?). Building upon the outcomes of the observations, we reflect upon the use of EHR data as the sole input for process mining and provide a glance to the future development of process mining in healthcare research.

This paper complements existing work reporting on the limitations of EHR data for process mining purposes by taking a different perspective. Existing works tend to consider the data as a starting point and, *e.g.*, try to understand the extent to which certain data quality issues appear in EHR data. This paper highlights another angle by taking one step back and starts from observing how nursing work is actually done, both the nursing activities and the accompanying data registration (if any). In this way, it provides a rich understanding of the context from which the EHR data will eventually emerge and which that data should sufficiently truthfully reflect.

The remainder of this paper is structured as follows. Section 2 outlines the related work. Section 3 describes the methodology of the observations. Section 4 summarizes the key results of the observations. Section 5 discusses the results and reflects upon them from a broader perspective. The paper ends with a conclusion in Section 6.

2. Related Work

In order to leverage EHR data to generate insights in care processes using process mining, various stages needs to be traversed. Figure 1 provides a high-level conceptualization of these stages. This section will highlight some pointers to the related work, structured along the lines of this conceptualization. Providing a full overview of literature is beyond the scope of this paper. For a more extensive overview on the state-of-the-art on process mining in healthcare, the reader is referred to review papers such as Rojas et al. (2016) and De Roock and Martin (2022).

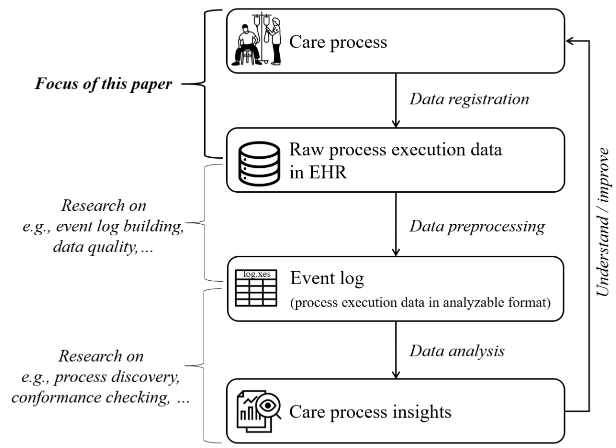

Figure 1. High-level conceptualization of the use of process mining with EHR data.

The predominant focus of existing research on the use of EHR data for process mining is *data analysis* — *i.e.*, the extraction of care process insights from an event log. This encompasses a wide variety of analyses such as process discovery, conformance checking, predictive process mining, and comparative process mining (van der Aalst, 2022) in various clinical application contexts. In the realm of process discovery, EHR data has, for instance, been used to discover the order of activities in the care pathway of patients receiving chemotherapy to treat breast cancer or colorectal cancer (Noshad et al., 2022), or the pathway of dental patients preceding teeth extraction under general anesthetic (Fox et al., 2022). An example of a paper that also incorporates conformance checking is Noshad et al. (2022), where EHR data is leveraged to discover the care pathway for emergency stroke patients and to assess the extent to which the trajectories of individual patients conform to the most common path. Predictive process mining using EHR data is, for instance, considered by Cremerius et al. (2022) in which process-related data is used to predict the discharge location of patients suffering from heart failure. Lastly, an illustration of comparative process mining is Yoo et al. (2016), where process mining on EHR data is used to compare the situation before and after new buildings in a hospital were taken into use. Papers on data analysis also highlight limitations in the available data such as process steps that do not leave a trail in the EHR or data entries that do not accurately reflect reality (Baker et al., 2017; Fox et al., 2022; Yoo et al., 2016).

An event log contains process execution data in analyzable format. To create an event log using raw process execution data from the EHR, extensive *data preprocessing* is typically required. This entails tasks such as event log building, data quality assessment, and data cleaning. While many of the approaches proposed in the broader process mining field can also be applied in the specific context of EHR data, some dedicated research has also been done. On the topic of event log building, Cremerius et al. (2023) propose a method to extract an event log from the MIMIC-IV database, a freely available EHR-based database from a US hospital. Other relevant contributions relate to the topic of data quality. In particular, Fox et al. (2018) present the Care Pathway Data Quality Framework (CP-DQF). This framework presents a stepwise approach to mark entries in an EHR-based event log as bad, compromised, or good. While bad events are unusable altogether, compromised events have issues but can still be used in particular settings. All data that is marked as good can be used for all analysis purposes. Another example is the work by Goel et al. (2023), which defines 6 digital health data imperfection patterns such as the fact that a data entry does not represent reality (*e.g.,* because the timestamp is incorrect). While Fox et al. (2018) and Goel et al. (2023) focus on assessing the event log that constitutes the input for process mining, Perimal-Lewis et al. (2016) propose using control-flow discovery to identify data quality problems. Abnormal flows in discovered control-flow models highlight potential areas in which timestamps might not be correctly recorded (Perimal-Lewis et al., 2016).

Data preprocessing methods take raw process execution data from the EHR as input. This data originates from registration actions that have taken place during the execution of a care process. Hence, *data registration* constitutes the key starting point for process mining. While the need to raise awareness regarding the importance of data registrations among staff members and to take initiatives to improve it has been articulated in literature (Martin et al., 2020; Martin et al., 2021; Zuidema-Tempel et al., 2022), dedicated research on this matter is currently largely lacking within the process mining field. Recently, the Odigos framework has been introduced and applied in a healthcare context to identify root causes of data quality issues in a process context (Andrews et al., 2022; Eden et al., 2023). Once these root causes are known, initiatives to tackle them can be taken in an effort to enhance data registration. The Odigos framework still considers a detected data quality issue as a starting point, instead of starting from understanding the care process. This paper takes this latter angle by observing how work is done in a care process and how this is registered in the EHR.

3. Methodology

In order to explore how accurately EHR data reflects the work done in a care process, an observation study has been conducted at a residential ward of a Belgian hospital. In this ward, patients are admitted for a particular period of time while they undergo treatment for one or more conditions. The nursing context is purposefully selected as nurses perform a wide variety of activities taking place during the daily operations of a ward. Examples of nursing activities include helping the patient to get dressed, changing the bed linen, administering medication, and measuring the patient's parameters.

The conducted observations focused on how nursing activities are being performed and how nurses register data regarding these activities in the EHR. The study took the form of non-participant observations — *i.e.*, observations in which the observer registers what happens without actively participating in the performed activities (Bougie & Sekaran, 2020). In practice, one specific nurse was followed during an entire shift. Additional shifts were observed until saturation was reached as further observations would not provide additional insights as no novel nursing activities nor data registration patterns would surface anymore (Flick, 2009). Each nurse could only be observed for one shift in the entire study to ensure that potential differences in behavior between nurses were sufficiently captured. Both morning and evening shifts were observed to account for differences in the nursing activities that might occur given the fact that morning and evening routines at a residential ward differ.

For each nursing activity performed during one of the observed shifts, the entries in the observation sheet summarized in Table 1 were recorded. Besides the *room* in which the activity took place, several aspects related to the execution of a nursing activity were noted: which activity it was (*activity*), the time at which it took place (*timestamp activity*), and details regarding its execution (*activity execution details*). These details relate, for instance, to whether the activity is performed by several nurses together, or whether a nurse multitasks and executes the activity together with another activity. Next to notes on the nursing activity itself, its registration in the EHR was also observed. This relates both to which information was recorded (*input EHR*) and the moment at which it was recorded (*timestamp EHR*). To accommodate recording additional information that might be relevant, the observation sheet also left room to enter additional observation notes (*additional notes*).

Table 1. Items observation sheet

Item	Brief description
Room	Room in which the activity took place
Activity	Description of the activity
Timestamp activity	Time at which the activity took place
Activity execution details	Notes on how an activity was executed (e.g. teamwork, multitasking, ...)
Input EHR	Description of the information recorded in the EHR
Timestamp EHR	Time at which the activity has been recorded in the EHR
Additional notes	Room to leave additional observation notes

4. Results

Consistent with the approach outlined in Section 3, additional shifts were observed until saturation was reached. This was the case after observing four shifts, with three nurses having an 8-hours shift and one nurse having a 7-hours shift. To confirm that saturation was actually reached, a fifth shift was observed half, which did not generate novel insights. This implies that a total of 35 working hours have been observed with 5 nurses.

To outline the key findings of the observations, a distinction is made between the observed nursing activities on the one hand (Section 4.1) and the observed data registration behavior of nurses on the other hand (Section 4.2).

4.1. Observed Nursing Activities

A total of 38 distinct nursing activities have been observed. Consistent with Gardner et al. (2010), a distinction is made between four key categories of nursing activities: direct care activities (*i.e.*, activities performed for the patient in their presence or their family's presence), indirect care activities (*i.e.*, activities performed on behalf of the patient, but not in their presence), service-related activities (*i.e.*, activities that are not specific to a patient), and personal activities (*i.e.*, activities that relate to the personal time of a nurse). Table 2 shows the number of distinct nursing activities observed for each of these categories, together with some examples. From the table, it follows that direct care (50.0%) and indirect care (39.50%) activities account for most of the observed nursing activities.

Table 2. Observed nursing activities

Category	N	Example(s)
Direct care	19	Administer medication, Measure patient's parameters, Wash patient, Move patient, ...
Indirect care	15	Collect required medical materials, Insert data in EHR, Brief the physician, Brief nurses from next shift, ...
Service-related	1	Compile reports
Personal	3	Have lunch, Make personal phone call, Talk with colleagues (social talk)

Besides the nursing activities that were performed, the observations also shed light on how work is organized. Four of the observed patterns are outlined below. Firstly, to determine which nurse performs an activity for a patient, the ward uses a patient assignment policy. This implies that patients are assigned to a nurse, who is responsible for all nursing activities related to the assigned patients (*i.e.*, integrated nursing). In the morning and evening shift, respectively 6 and 12 patients are assigned to each nurse.

Secondly, batching behavior takes place when nurses purposefully organize the execution of an activity in such a way that they can perform that activity for multiple patients in a consolidated time frame (Martin et al., 2017). Batching has, for instance, been observed when medication is being prepared as nurses perform this activity for multiple patients in one time frame. Similarly, batching occurred when administering medication as this was part of a dedicated tour in which a nurse sequentially visited several patients.

Thirdly, nurses regularly multitask, implying that multiple activities are performed for a particular patient during a single visit. Note that this differs from batching: while batching implies that a nurse performs one activity for several patients in a consolidated time frame, multitasking involves the execution of different activities for one specific patient during one visit. When considering multitasking, a distinction can be made between task switching and dual tasking. While task switching refers to behavior in which a nurse alternates between several different activities, dual tasking reflects several activities being performed simultaneously (Appelbaum et al., 2008). Task switching has, for instance, been observed when a nurse guides a patient to the bathroom, tidies the patient's bed, and guides him/her back to the bed. An example of dual tasking has been observed when the blood pressure and the body temperature of a patient were taken.

Finally, it is observed that nurses can opt to perform certain activities in team — *i.e.*, together with one or more colleagues. This happens, for instance, when the assistance of a colleague is needed to move obese patients from their bed to the chair. At other times, nurses prefer to work together with a colleague and jointly visit all patients that one of them is responsible for in a particular tour across the ward.

4.2. Observed Data Registration Behavior

From the observations of the data registration behavior, it follows that only 4 out of the 38 observed activities leave an explicit trail in the EHR: the registration of the patient's parameters (*e.g.*, blood pressure and body temperature), the preparation of the medication for a patient from the medicine cabinet, the administration of medication, and a change of the patient's position to prevent bedsores (in the bed, or between the chair and the bed). For some of the other activities, indirect indications of their occurrence could be derived from the free text notes that a nurse can add to the patient's file. For instance: when notes are added on a patient's state of mind, this indicates that a conversation between nurse and patient took place.

When considering the time at which an activity is registered in the EHR vis-à-vis its actual execution, it is observed that EHR registration rarely takes place at the moment at which an activity is actually performed. Most often, activities are recorded in the EHR some time *after* their execution. Work at the ward is typically organized in tours, in which a nurse performs one or more activities for the patients for which he/she is responsible. What is frequently observed is that a nurse performs a tour and makes entries in the EHR some time after the tour has ended. By default, EHR entries are marked with the timestamp at which the entry is made in the system. However, nurses can also enter the timestamp at which the activity has been performed when the registration time does not correspond with the execution time. It has been observed that when nurses use this option, they enter the indicative time at which the activity took place (*e.g.,* 5.00 p.m.). This time tends to be the same for all activities performed in a particular tour.

While the data registration pattern in which an activity is recorded some time after being performed is the most common, activities were sometimes also recorded in the EHR *before* they were executed. For instance: when preparing the medication trolley for the medication tour, a nurse might already record the administration of the medication to the patients in the EHR. The actual administration takes place at a later moment as the medication tour starts some time after the trolley has been prepared.

Only in exceptional cases, EHR entries were made immediately when the activity was performed. Nevertheless, nurses indicated that they saw the advantage of immediately recording all relevant information in the EHR as this would lead to more accurate and complete information, as well as mitigate the risk of forgetting to register something. However, the time pressure that they experience in parts of their shift led them to defer data registration to a calmer period.

An additional factor to take into account is the patient assignment policy that the ward applies (as introduced in Section 4.1), which also implies that only the responsible nurse can make recordings in the EHR for a particular patient. In case another nurse takes over a particular activity for that patient, all relevant information needs to be transferred (*e.g.*, on a piece of paper) such that the responsible nurse can enter it into the EHR. Besides being error-prone, this transfer of information also causes a delay in the registration of data in the EHR. It has, for instance, been observed that a colleague made some notes regarding patients on paper, after which the responsible nurse inserted the information in the EPD 16 minutes later.

With respect to the other work organisation patterns highlighted in Section 4.1, no explicit EHR recordings have been observed. Firstly, when batching took place, data was recorded at another point in time. Even when the nurse altered the timestamp when making entries in the system, these were typically all set to the same timestamp. While having the same activity sharing an identical timestamp over multiple patients might be a preliminary indication that batching took place, the EHR data will not provide any further pointers to figure this out. Secondly, multitasking will not be retrievable from the EHR for the observed shifts because at most one activity led to a registration in the EHR. Finally, teamwork will not become visible in the EHR as only the responsible nurse can record information in the EHR. Hence, it will not be reflected whether an activity has been performed by multiple nurses. It is even not guaranteed that the responsible nurse actually executed the activity as nurses might help each other out.

5. Discussion

The discussion section has a threefold focus. Firstly, the use of EHR data as a basis for process mining is reflected upon (Section 5.1). Secondly, a glance to the future development of process mining in healthcare research is provided (Section 5.2). Finally, the contributions are summarized and the limitations of the paper are recognized (Section 5.3).

5.1. Reflections on EHR data as a basis for process mining

The observations of both the performed nursing activities as well as the EHR data registration behavior provide a basis to reflect upon the extent to which EHR data is a truthful basis for process mining. In general, the results show that EHR data will provide a highly fragmented and inaccurate view of how nursing work gets done at the observed ward. Foremost, many nursing activities are not recorded in the EHR at all, implying that they will remain invisible when EHR data is used as a basis to study the process. Even when an activity leads to a registration in the EHR, the timestamp often does not correspond to the time at which the activity actually took place. Moreover, the order in which activities are recorded could even differ from the order in which they were actually performed. Besides the potential inaccuracy of the timestamp, it has also been observed that hardly any contextual information about the activity execution is recorded in the EHR. For instance: it will remain hidden whether a nurse performed this activity individually or whether teamwork took place.

To place these observations into context, it is important to recall that an EHR system is developed to support the operations of a hospital by providing a system to capture the health-related information of a patient (Häyrinen et al., 2008). Hence, it is not designed as a system to track nursing activities throughout the patient's stay at a ward. Neither is such a system put in place for the purposes of conducting data analyses. As the implementation of an EHR system entails significant investments from a hospital both in terms of software as well as to train staff to use the system (Adler-Milstein et al., 2015), these systems are likely to remain in place for extended periods of time. Hence, it constitutes a reality that the process mining community should be aware of.

The aforementioned findings regarding EHR data present important limitations for its use as the sole starting point for process mining. To conduct process mining, an event log should minimally contain an ordered overview of the relevant activities executed for each case (*e.g.*, for each patient), where this order is typically expressed by means of a timestamp (van der Aalst, 2016). If only a small fraction of the activities that actually take place are recorded and the associated timestamps often provide a distorted view of the actual timings, one might question the extent to which process mining can deliver upon its promise of increasing transparency in end-to-end processes (Martin et al., 2020).

Of course, data requirements will depend upon the analysis question(s) at hand. If a hospital, for instance, wants to understand in which order various medications are administered to the patient during his/her stay, it could be sufficient if the EHR reflects which medication was provided to a patient with a time accuracy of an hour. Data on other activities such as helping the patient to get dressed and delivering a meal might be irrelevant. In contrast, when a hospital wants to leverage process mining to understand the workload and work organization of nurses, a full and accurate coverage of the activities that are performed and who performed them is needed. Consequently, the extent to which EHR data covers the executed activities, as well as the accuracy and richness (in terms of contextual variables) of the recorded data will determine the versatility of the analysis questions for which process mining can provide reliable answers.

If process mining has the ambition to become a key ingredient in establishing a process-aware and data-driven culture within hospitals, it should be able to reliably answer a great variety of analysis questions. The fact that EHR data is currently unlikely to be a solid basis to provide these answers can be detrimental to the uptake of process mining in healthcare. If the delivered insights are not perceived as accurate (due to the way in which data is recorded) or useful (given its limited coverage of the activities in a care process), hospitals' belief in process mining as a means to improve care processes in a data-driven way can quickly reduce. In the longer run, this can impede the further development of this promising domain in which exchanges between researchers and hospitals are crucial (Martin et al., 2020).

One could argue that it is the responsibility of hospitals to ensure that the required data with a sufficient quality level is available if they want to benefit from what process mining has to offer. However, this attitude would contradict the ambition of the process mining in healthcare community to have societal impact, which has been vocalized in Munoz-Gama et al. (2022). Moreover, moving this responsibility fully to hospitals is also not in the interest of the research community as the current limitations of EHR data might reduce the credibility of process mining in the longer run, as highlighted above. To move hospitals towards actions to critically assess and improve the gathering of process-related data, thorough awareness of the existence of process mining and its potential to improve their operations are necessary conditions. It is currently unclear to what extent awareness on these matters exists within hospitals, but observations such as the rather limited uptake of process mining in hospitals outside a research context suggests that further efforts are still needed (Martin et al., 2020). The process mining in healthcare research community should play an active role here.

From the previous, it follows that EHR data by itself might not constitute a solid basis to progress towards the systematic use of process mining in hospitals. The conducted observations demonstrate some of the key limitations such as the limited coverage of the actual work that is done in a process, and the mismatch between activity execution time and data registration time. Such limitations will influence the impact that process mining can generate and, in the longer run, the credibility of process mining in healthcare. Process mining in healthcare is a maturing discipline that is gradually approaching a crossroads: either it convinces hospitals to start embedding it in their daily operations, or it leads to disillusion because the generated process insights gathered from the available EHR data are too fragmented or inaccurate. Consequently, in the years to come, the research community should invest efforts in clearly demonstrating the full potential of process mining in healthcare, as well as in providing novel methods and tools that facilitate hospitals to leverage this potential by targeting the input side of process mining. The next subsection provides a more elaborate glance to the future against the background of the reflections raised in this subsection.

5.2. Glance to the Future

The prior subsection reflected upon the findings of the observations and their broader implications. This clearly highlights that the process mining in healthcare community should not only focus on the development of novel process mining methods to analyze EHR data, but also invest in better understanding and strengthening the foundations of the data that is used as input. In that respect, three key considerations to inspire future research can be distinguished.

Firstly, the process mining in healthcare community should invest more in *understanding care processes and nursing work*. At the moment, the data is often considered the starting point and limited attention is attributed to sufficiently grasp the characteristics of the care process itself and how nurses get things done. However, having such an understanding is a prerequisite to correctly interpret and use the data that the process generates. The observations have demonstrated that raw EHR data can present a biased view of the care process. When inadequate attention is attributed to understanding the process and nursing work, such biases can easily be propagated into the remainder of the analysis, potentially leading to inaccurate or even misleading results, which can be detrimental for the support of process mining. This also underlines the critical importance of conducting research projects in close collaboration with healthcare professionals.

Secondly, there is a need to *increase the transparency of how accurate the EHR data represents the care process*. This information is highly relevant as it will influence the reliability of process mining output. Hence, a process mining study should report upon the prevailing discrepancies between the care process and the recorded EHR data, as well as explicitly reflect upon their implications (*e.g.*, the research questions that can no longer be answered in a reliable way). In this realm, measures can also be developed to, for instance, express the degree to which EHR data covers the key activities in the care process. To determine the value of such a measure, observations of the care process or structured discussions with domain experts would be needed.

Finally, research on *supporting the improvement of data registration in care processes* is warranted. Improved data registration relates to both ensuring that process execution data more correctly represents the care process (*e.g.,* more accurate timestamps), as well as enriching the data such that more contextual information is included (*e.g.,* to make collaboration between nurses visible). This can be achieved by altering the way in which the EHR is used in a care process, *e.g.*, by requiring that activities are immediately recorded when they are executed, by extending the nursing activities that need to be registered, and by adding additional fields that need to be entered. However, making nurses record a multitude of additional activities and values surpasses the primary goal of the EHR, *i.e.*, capturing a patient's health-related information (Häyrinen et al., 2008). Moreover, it would place a significant additional administrative burden upon nurses, who are currently already confronted with a high workload (Van den Heede et al., 2023). As a consequence, other solutions should be explored. A promising direction includes the use of mobile apps in which a nurse can quickly register activities with a few clicks. While the use of an app still requires action by a nurse, indoor location systems can also be used. These systems collect information about the location of a nurse at a particular point in time (Bendavid, 2016), which can be leveraged to identify the execution of nursing activities.

When exploring ways to improve data registration, a trade-off should always be made between the benefits in terms of enhanced data registration and the associated analysis potential on the one hand, and the costs that will be incurred on the other hand. Costs should be broadly interpreted as it encompasses not only the investment in hardware or software, but also the burden that is placed on nurses and the impact that, *e.g.*, indoor location tracking can have on employee well-being. There is ample room for future research to investigate such trade-offs.

Improving data registration will require investments from hospitals as even merely freeing up precious time from a team already has a cost. In order to justify such investments, hospitals need to be convinced of process mining's value to improve their operations. At the same time, the observed limitations of the data captured by the EHR systems make it particularly difficult to fully demonstrate what process mining can offer to hospitals. A way out of this apparent chicken-and-egg situation is to set up pilot studies in which EHR data is corrected and enriched with additional sources of process execution data. These sources can include technologies such as mobile apps or indoor location systems as highlighted above, but also observation studies. The goal of these studies should be to clearly demonstrate what process mining can offer to hospitals and how it can actually contribute to the improvement of care processes. Such studies will be pivotal in providing evidence that investing in improved data registration is worthwhile. To guide the focus of these research efforts, insights can be gathered on the awareness about the existence of process mining in hospitals, as well as the potential that they see. Until now, no systematic understanding has been gathered on these matters.

5.3. Contributions and Limitations

To frame this paper's contributions, it should be acknowledged that some of the findings that emerged from the observations have already been coined in prior literature. For example: Rule et al. (2020) indicate that not all activities give rise to a registration in the EHR and, hence, that EHR data does not provide a full view of the work being done. Similarly, the mismatch between the activity execution time and the timestamp in recorded data has been highlighted in, among others, Vanbrabant et al. (2019) and Martin and Bergs (2020). This paper confirms these findings based on an observation study and provides a significantly richer contextual view of how work is performed in reality than prior literature. Moreover, this paper reflects more extensively on the findings and their broader implications for process mining in healthcare. In this way, it extends the knowledge base by focusing on how truthful EHR data reflects the execution of a care process, a topic that receives very little attention despite its crucial impact on the potential of EHR data as a basis for process mining.

The outcomes of this study have to be reflected against its limitations. Firstly, the observation study has been conducted in one setting, *i.e.* one ward in a hospital. Within this setting, sufficient observations were made as saturation was reached. However, we cannot formally claim that the findings also hold in other healthcare settings. Nonetheless, the fact that some findings corroborate views already articulated in literature can be seen as an indication that the results are likely to, at least partly, hold at a more general level as well. Secondly, the Hawthorne effect might occur within the context of direct observations, reflecting the behavioral change that can arise because subjects know they are being observed (Brysbaert, 2006). The fact that consistent behavioral patterns were observed across nurses over different shifts support the reliability of the findings that emerged.

6. Conclusion

This paper presents the outcomes of an observation study at a Belgian hospital in which the activities that nurses perform have been observed, as well as how they register data in the EHR. The conducted observations highlight that EHR data will provide a highly fragmented and inaccurate view of how work gets done as, *e.g.*, many activities are not recorded at all and, for others, the recorded timestamps do not correspond to the time at which the activity is actually executed. These findings constitute important limitations for the use of EHR data as the sole starting point for process mining. As a consequence, the research community should attribute more attention to establishing research initiatives around data registration to better understand and strengthen the foundations of the input data.

To further substantiate the findings that emerged from the observation study, a large-scale study can be set up in which nurses are observed in a wide variety of healthcare settings. This would enable us to gain insights into the commonalities and differences between different healthcare contexts, as well as their implications on the use of process mining in these settings. Various other directions for future research have already been highlighted in Section 5.2.

References

Adler-Milstein, J., DesRoches, C. M., Kralovec, P., Foster, G., Worzala, C., Charles, D., Searcy, T., & Jha, A. K. (2015). Electronic health record adoption in US hospitals: progress continues, but challenges persist. *Health Affairs*, *34*(12), 2174–2180.

Andrews, R., Emamjome, F., ter Hofstede, A. H., & Reijers, H. A. (2022). Root-cause analysis of process-data quality problems. *Journal of Business Analytics*, *5*(1), 51–75.

Appelbaum, S. H., Marchionni, A., & Fernandez, A. (2008). The multi-tasking paradox: perceptions, problems and strategies. *Management Decision*, *46*(9), 1313–1325.

Baker, K., Dunwoodie, E., Jones, R. G., Newsham, A., Johnson, O., Price, C. P., Wolstenholme, J., Leal, J., McGinley, P., Twelves, C., et al. (2017). Process mining routinely collected electronic health records to define real-life clinical pathways during chemotherapy. *International Journal of Medical Informatics*, *103*, 32–41.

Bendavid, Y. (2016). RTLS in hospitals: technologies and applications. In *Encyclopedia of e-commerce development, implementation, and management* (pp. 1868–1883). IGI Global.

Bougie, R., & Sekaran, U. (2020). *Research methods for business*. Wiley.

Brysbaert, M. (2006). *Psychologie*. Academia Press.

Cremerius, J., König, M., Warmuth, C., & Weske, M. (2022). Patient discharge classification based on the hospital treatment process. *Lecture Notes in Business Information Processing*, *433*, 314–326.

Cremerius, J., Pufahl, L., Klessascheck, F., & Weske, M. (2023). Event log generation in MIMIC-IV. *Lecture Notes in Business Information Processing*, *468*, 302–314.

De Roock, E., & Martin, N. (2022). Process mining in healthcare – an updated perspective on the state of the art. *Journal of Biomedical Informatics*, *127*, paper 103995.

Eden, R., Syed, R., Makasi, T., Andrews, R., ter Hofstede, A. H. M., Wynn, M. T., Donovan, R., Eley, R., & Staib, A. (2023). Revealing the root causes of digital health data quality issues: a qualitative Investigation of the Odigos framework. *Proceedings of the 56th Hawaii International Conference on System Sciences*, 3035–3044.

Flick, U. (2009). *An introduction to qualitative research*. Sage.

Fox, F., Whelton, H., Johnson, O. A., & Aggarwal, V. R. (2022). Dental extractions under general anesthesia: new insights from process mining. *JDR Clinical & Translational Research*.

Fox, F., Aggarwal, V. R., Whelton, H., & Johnson, O. (2018). A data quality framework for process mining of electronic health record data. *Proceedings of the 2018 IEEE International Conference on Healthcare Informatics*, 12–21.

Gardner, G., Gardner, A., Middleton, S., Della, P., Kain, V., & Doubrovsky, A. (2010). The work of nurse practitioners. *Journal of Advanced Nursing*, *66*(10), 2160–2169.

Goel, K., Sadeghianasl, S., Andrews, R., Ter Hofstede, A., Wynn, M. T., Kapugama Geeganage, D., Leemans, S. J., McGree, J., Eden, R., Staib, A., et al. (2023). Digital health data imperfection patterns and their manifestations in an Australian digital hospital. *Proceedings of the 56th Hawaii International Conference on System Sciences*, 3235–3244.

Häyrinen, K., Saranto, K., & Nykänen, P. (2008). Definition, structure, content, use and impacts of electronic health records: a review of the research literature. *International Journal of Medical Informatics*, *77*(5), 291–304.

Kusuma, G. P., Kurniati, A. P., Rojas, E., McInerney, C. D., Gale, C. P., & Johnson, O. A. (2021). Process mining of disease trajectories: a literature review. *Proceedings of the 2021 Medical Informatics Europe Conference*, 457–461.

Mans, R. S., van der Aalst, W. M. P., & Vanwersch, R. J. B. (2015). *Process mining in healthcare: evaluating and exploiting operational healthcare processes*. Springer.

Martin, N. (2021). Data quality in process mining. In C. Fernandez-Llatas (Ed.), *Interactive process mining in healthcare* (pp. 53–79). Springer.

Martin, N., & Bergs, J. (2020). Patient flow data registration: a key barrier to the data-driven and proactive management of an emergency department. *International Emergency Nursing*, *53*, paper 100932.

Martin, N., De Weerdt, J., Fernández-Llatas, C., Gal, A., Gatta, R., Ibáñez, G., Johnson, O., Mannhardt, F., Marco-Ruiz, L., Mertens, S., et al. (2020). Recommendations for enhancing the usability and understandability of process mining in healthcare. *Artificial Intelligence in Medicine*, *109*, paper 101962.

Martin, N., Fischer, D. A., Kerpedzhiev, G. D., Goel, K., Leemans, S. J., Röglinger, M., van der Aalst, W. M., Dumas, M., La Rosa, M., & Wynn, M. T. (2021). Opportunities and challenges for process mining in organizations: results of a Delphi study. *Business & Information Systems Engineering*, *63*, 511–527.

Martin, N., Swennen, M., Depaire, B., Jans, M., Caris, A., & Vanhoof, K. (2017). Retrieving batch organisation of work insights from event logs. *Decision Support Systems*, *100*, 119–128.

Martin, N., Van Houdt, G., & Janssenswillen, G. (2022). DaQAPO: supporting flexible and fine-grained event log quality assessment. *Expert Systems with Applications*, *191*, 116274.

Munoz-Gama, J., Martin, N., Fernández-Llatas, C., Johnson, O. A., Sepúlveda, M., Helm, E., Galvez-Yanjari, V., Rojas, E., Martinez-Millana, A., Aloini, D., Amantea, I. A., Andrews, R., Arias, M., Beerepoot, I., Benevento, E., Burattin, A., Capurro, D., Carmona, J., Comuzzi, M., … Zerbato, F. (2022). Process mining for healthcare: characteristics and challenges. *Journal of Biomedical Informatics*, *127*, paper 103994.

Noshad, M., Rose, C. C., & Chen, J. H. (2022). Signal from the noise: a mixed graphical and quantitative process mining approach to evaluate care pathways applied to emergency stroke care. *Journal of Biomedical Informatics*, *127*, paper 104004.

Perimal-Lewis, L., Teubner, D., Hakendorf, P., & Horwood, C. (2016). Application of process mining to assess the data quality of routinely collected time-based performance data sourced from electronic health records by validating process conformance. *Health Informatics Journal*, *22*(4), 1017–1029.

Poba-Nzaou, P., & Uwizeyemungu, S. (2019). Variation in electronic health record adoption in European public hospitals: a configurational analysis of key functionalities. *Health and Technology*, *9*, 439–448.

Rojas, E., Munoz-Gama, J., Sepúlveda, M., & Capurro, D. (2016). Process mining in healthcare: a literature review. *Journal of Biomedical Informatics*, *61*, 224–236.

Rule, A., Chiang, M. F., & Hribar, M. R. (2020). Using electronic health record audit logs to study clinical activity: a systematic review of aims, measures, and methods. *Journal of the American Medical Informatics Association*, *27*(3), 480–490.

Van den Heede, K., Balcaen, K., Bouckaert, N., Bruyneel, L., Cornelis, J., Sermeus, W., & Van de Voorde, C. (2023). Improving hospital nurse staffing during the pandemic: implementation of the 2019 fund for healthcCare staff in Belgium. *Health Policy*, *128*, 69–74.

van der Aalst, W. M. P. (2016). *Process mining: in action*. Springer.

van der Aalst, W. M. P. (2022). Process mining: a 360 degree overview. In W. M. P. van der Aalst & J. Carmona (Eds.), *Process Mining Handbook* (pp. 3–34). Springer.

Vanbrabant, L., Martin, N., Ramaekers, K., & Braekers, K. (2019). Quality of input data in emergency department simulations: framework and assessment techniques. *Simulation Modelling Practice and Theory*, *91*, 83–101.

Yoo, S., Cho, M., Kim, E., Kim, S., Sim, Y., Yoo, D., Hwang, H., & Song, M. (2016). Assessment of hospital processes using a process mining technique: outpatient process analysis at a tertiary hospital. *International Journal of Medical Informatics*, *88*, 34–43.

Zerbino, P., Stefanini, A., & Aloini, D. (2021). Process science in action: a literature review on process mining in business management. *Technological Forecasting and Social Change*, *172*, paper 121021.

Zuidema-Tempel, E., Effing, R., & van Hillegersberg, J. (2022). Bridging the gap between process mining methodologies and process mining practices. *Lecture Notes in Business Information Processing*, *458*, 70–86.