

6-14-2021

DEUCE: A Methodology for Detecting Unauthorized Access of Electronic Health Records Using Process Mining

Victor van Anandel
NFIR B.V., v.vanandel@nfir.nl

Iris Beerepoot
Utrecht University, i.m.beerepoot@uu.nl

Xixi Lu
Utrecht University, x.lu@uu.nl

Inge van de Weerd
Utrecht University, i.vandeweerd@uu.nl

Hajo A. Reijers
Utrecht University, h.a.reijers@uu.nl

Follow this and additional works at: https://aisel.aisnet.org/ecis2021_rp

Recommended Citation

van Anandel, Victor; Beerepoot, Iris; Lu, Xixi; van de Weerd, Inge; and Reijers, Hajo A., "DEUCE: A Methodology for Detecting Unauthorized Access of Electronic Health Records Using Process Mining" (2021). *ECIS 2021 Research Papers*. 66.

https://aisel.aisnet.org/ecis2021_rp/66

This material is brought to you by the ECIS 2021 Proceedings at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2021 Research Papers by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

DEUCE: A METHODOLOGY FOR DETECTING UNAUTHORIZED ACCESS OF ELECTRONIC HEALTH RECORDS USING PROCESS MINING

Research Paper

Victor van Andel, NFIR B.V., The Hague, Netherlands

Iris Beerepoot, Utrecht University, Utrecht, Netherlands, i.m.beerepoot@uu.nl

Xixi Lu, Utrecht University, Utrecht, Netherlands

Inge van de Weerd, Utrecht University, Utrecht, Netherlands

Hajo Reijers, Utrecht University, Utrecht, Netherlands

Abstract

Hospitals have a keen interest to root out unauthorized access of Electronic Health Records (EHRs). The retrospective analysis of suspicious EHR views may help to take preventive measures against such access. However, investigating EHR access manually is labor-intensive and only allows for checking a small sample of cases. We explore how process mining techniques can be used to support the detection of unauthorized views. In the context of EHR access, it is easier to define authorized views with certainty than to detect unauthorized views. Therefore, we propose DEUCE: a methodology that focuses on identifying authorized behavior, such that unauthorized views are distinguished, and can be investigated more in-depth. We evaluate the methodology in the form of a case study at a Dutch hospital. As a result of this study, the hospital has adopted the approach in favor of their traditional approach and unauthorized EHR access is now more efficiently detected.

Keywords: Process Mining, Electronic Health Records, Unauthorized Access, Security, Deviations.

1 Introduction

The emergence of Health Information Systems has enabled the digital availability of information that is needed by health workers to accomplish their job. Information about a patient is stored in their Electronic Health Record (EHR), which has become an integral part of modern healthcare (Borycki et al., 2011). With the introduction of the General Data Protection Regulation (GDPR), hospitals need to satisfy strict rules with regards to EHRs and privacy of patients. In order to enforce the rules, elaborate security and privacy requirements have been formulated to protect the sensitive data that they contain (Rodrigues et al., 2013). Deviations from security and privacy requirements can result in major fines for the organisations involved. Therefore, it is vital for hospitals to avoid unauthorized access of the EHR by employees. However, distinguishing unauthorized from authorized access is not an easy task. Consider the following example:

A male patient is brought into the emergency department by ambulance. The doctor assigned to this patient quickly checks his EHR for his medical history. Formally, the patient's visit is not yet registered as such in the system. GDPR law prescribes that only healthcare professionals with a formal treatment relationship with the patient are authorized to view his record. As this is an emergency situation, the doctor utilizes the break the glass mechanism of the system in order to view the file. In this situation, the patient's file has been viewed for a legitimate reason. Consider, by contrast, a doctor from another department hearing that an old acquaintance has been admitted to the

emergency department. Out of interest, this doctor views the patient's file and uses the break the glass mechanism to do so. As this doctor does not have and will not have a treatment relation with the patient, this is an example of unauthorized EHR viewing .

Although it is important for hospitals to monitor EHR access and take preventive measures when unauthorized access is detected, doing so manually is highly labor-intensive. Moreover, access rules are flexible to account for emergencies, making the specification of authorized access difficult. In the example above, a privacy officer would need to check whether the doctor became the patient's attending doctor after using the break the glass mechanism: e.g. by checking whether the doctor operated on the patient or prescribed medication. In practice, this means privacy officers are only able to check a sample of all cases by retrospectively checking access. Fortunately, the use of data mining techniques, in particular process mining, can offer a solution to this problem. Process mining techniques use so-called event logs, extracted from an IT system, to perform process analyses on those data (van der Aalst, 2016). It allows for analyzing what has happened when, but also what happened before and after. Therefore, it can provide sequential information, which other data mining techniques cannot. Thus, process mining is of particular value for distinguishing authorized from unauthorized access for examples such as the one mentioned. In addition, process mining techniques can provide insights into the resources performing tasks, time aspects involved, and more (Mannhardt et al., 2016). In situations other than the one in the example, these aspects may turn out to be important. Using process mining to distinguish authorized from unauthorized viewing by healthcare professionals might allow for monitoring a larger set of cases and detect more instances of unauthorized views. Therefore, we investigate the question: *how can process mining be used to detect unauthorized EHR access?*

In this paper, we adapt the well-established Process Mining Project Methodology PM² (van Eck et al., 2015), in order to propose a methodology for using process mining to detect unauthorized EHR access. Specifically, we adopt the six steps from the original methodology, the iterative cycles, and the involvement of business experts during different steps in the project. On top of that, we provide detailed support for our specific goal of iteratively detecting unauthorized access. We applied and evaluated the proposed approach within a case study at a Dutch hospital. Our main scientific contribution is a step-by-step approach that can be used by hospitals to investigate EHR access. It has the potential to be used to more generally detect deviations of information systems users. Our proposal includes an introduction of two metrics to evaluate the extent to which the results of the analysis are precise. Unique to our methodology is the focus on excluding authorized behavior to arrive at a set of unauthorized cases, by defining green flags. Practically, our methodology provides an improvement in terms of time and coverage compared to the traditional approach of hospitals to manually analyze small samples of cases of EHR access. As a result of this study, the hospital involved adopted the approach to more efficiently detect unauthorized EHR access.

2 Background Literature

2.1 Privacy and Legislation

Hospitals are subject to several information security regulations, which differ from one region to the other. In the European Union, for example, hospitals must adhere to the European General Data Protection Regulation (GDPR), which is concerned with the processing of personal data of all citizens of the European Union. In addition, countries also have more specific legislation related to information security. Apart from regional differences, there are also differences in which rules apply to certain parts of information systems. For example, different rules might hold for different parts of the EHR. Murphy et al. (2014) distinguish between three types of patient information: protected health information (anything about a patient's medical status), personally identifiable information (anything that could identify the patient), and generalized information (anything else). A different level of discretion may be required depending on the kind of information that is dealt with. This could be embodied by restricting certain parts of the records for certain kinds of employees.

The most common form of system access control in hospitals is *user authentication* (Murphy et al., 2014). By requiring users to be logged in, every action that is performed within the system can be traced back to a user. Even when a malicious action is not prevented, it can be discovered, and the offender can be held accountable. Simply knowing that every action can be traced back to the person might prevent people from acting inappropriately.

Another topic tied to access control is determining which permissions should be given to which user. This is referred to as *user provisioning*. Give a user too many permissions and questions will be asked. Give too little and the user will not be able to do his job. User provisioning is quite complicated in reality, especially in healthcare where information needs often shift (Koppel et al., 2015). Health workers can have multiple roles; they may even move between different hospitals that use the same system. Patients can also move from one department to the other or be treated at multiple ones simultaneously. In situations such as the example in the introduction, it may save lives to quickly gain access to a patient's record. Therefore, some hospitals employ a break the glass mechanism, which health workers can use to transcend their permissions and gain access to a specific record in case of an emergency (Lovis et al., 2007).

2.2 Process Mining, Compliance, and Conformance Checking

Process mining is a collection of techniques that bridges the gap between traditional model-based process analysis and data-centric analysis techniques (van der Aalst, 2016). The application of process mining on healthcare processes has been well documented (Mans et al., 2015; Rojas et al., 2016). The healthcare sector is known for complex, ad-hoc, and multidisciplinary processes, posing particular challenges in process analysis endeavours (Homayounfar, 2012). Using process mining to provide insights into bottlenecks in the process and pinpointing opportunities for improvement can bring about direct improvements in patient care. For process mining techniques to work in healthcare, as well as in any other sector, an *event log* is required. An event log consists of a number of *events*. An event always contains an *activity*, e.g. a patient arriving at the emergency department or a doctor prescribing medication. Other information required in the event log is the *case* the event refers to, e.g. the patient involved, and the *timestamp*, i.e. the exact date and time the event occurred. With this information the process can be analyzed. Additional data such as the resources performing the activities or other attributes, if available, can enrich the analysis results.

Within process mining, three families of techniques can be distinguished (van der Aalst, 2016). The first relates to the *discovery* of process models without any a-priori knowledge. The second family of techniques is focused on *conformance* of what happened in reality as opposed to what was expected based on the process model. The third is related to *enhancement*, meaning that existing process models are enriched with information from the logs. In the context of this study, we focus especially on the family of conformance techniques, as we are interested in checking EHR access in relation to legislation. Commonly-used techniques within conformance checking are *rule checking*, *token replay*, and *alignments* (Carmona et al., 2018). A related field is that of *compliance checking*, which has largely focused on process modelling but has also recently explored approaches for discovering deviations (e.g. Ramezani et al., 2012). However, the latter has mostly looked at checking log behavior against predefined rules, while this does not solve our problem of defining such rules. In Chapter 2, we will describe how we incorporate these existing ideas towards the goal of checking EHR access.

2.3 Using a Red Flag Approach

For conformance checking techniques to work, there needs to be some normative model against which the log can be checked. Let us return to our example of unauthorized use of the break the glass mechanism. If we want to distinguish unauthorized from authorized use, we would need to specify exactly in which situations the doctor is authorized to access a patient file, and in which situations this is unauthorized. As was mentioned in section 2.1, access control is particularly complex in healthcare. For example, should we categorize all situations in which a doctor did not operate on a patient after

using the break the glass mechanism as unauthorized EHR access, the result would be a large number of false positives (Baader and Krcmar, 2018). Should we try to capture all exceptional situations in one statement, it would become uncomprehensibly long and complex. However, when presented with a specific example of a doctor using the break the glass mechanism and consequently treating that same patient, it is easy to determine this as authorized. Therefore, in situations such as described here, it is easier to say with certainty that some behavior that we see in the data is an example of authorized access, than it is to specify unauthorized access.

Drawing on principles of a so-called red flag approach, we aim to provide support for systematically distinguishing unauthorized from authorized access of systems. The red flag approach is often applied in the accounting field to detect fraud (Albrecht et al., 2011). The core idea of the approach is that unusual behavior is flagged and consequently investigated in more detail. Important to note is that at this point, the behavior is not yet categorized as fraud, but studied more in-depth. The involvement of a knowledgeable investigator is therefore required. Additionally, the approach assumes that there is no need to further investigate behavior that is not flagged. By doing so, it allows one to zoom in on suspicious behavior and not spend time on investigating normal behavior. Baader and Krcmar (2018) have shown how the red flag approach can be combined with process mining and can significantly bring down the number of false positives. To the best of our knowledge, theirs is the only work combining the two approaches. The main idea we pursue is to combine process mining and the red flag approach to filter out instances of unauthorized EHR access. However, as adopting the red flag approach as-is still leaves us with the problem of specifying unauthorized access, we deviate from the traditional approach. In the following section, we propose DEUCE: a methodology for detecting unauthorized viewing, which includes support for specifying unauthorized access.

3 DEUCE: A Methodology for Detecting Unauthorized Access of Electronic Health Records

In this section, we describe our methodology based on process mining for the detection of unauthorized EHR access, drawing from the previously mentioned ideas behind conformance checking and the red flag approach. Taking the break the glass example, we will flag every situation where the mechanism is used as suspicious. Since it is difficult to capture all situations in which the mechanism is unauthorized, we aim to find those situations in which it was used justly, i.e. we assign *green flags*. These instances do not need to be further investigated. In other words, we focus on extracting authorized behavior, such that we are left with unauthorized behavior. The list of green flags is continuously expanded to increase the precision of the detection.

3.1 Process Mining Methodologies

In recent years, several process mining methodologies have been introduced (e.g. van Eck et al., 2015; Zerbino et al., 2018), some even specific to analyzing healthcare processes (Rebuge and Ferreira, 2012; Rojas et al., 2017). The methodology we propose is an adaptation of the well-established Process Mining Project Methodology (PM²) by Van Eck et al. (2015). There are three main reasons for choosing this methodology as a basis for detecting unauthorized EHR access. First, it includes a step in which conformance checking can be applied. Second, iterative analysis cycles are inherent to the methodology. Third, it is the only methodology that explicitly specifies the types of experts involved in each step.

PM² is a general methodology that applies to a very diverse application of process mining techniques. We adopt the six steps from the original methodology, the iterative cycles, and the involvement of business experts during different steps in the project. On top of that, we provide detailed support for our specific goal of iteratively detecting and refining green flags. The proposed methodology is visualized in **Figure 1**. We will describe each of the steps below.

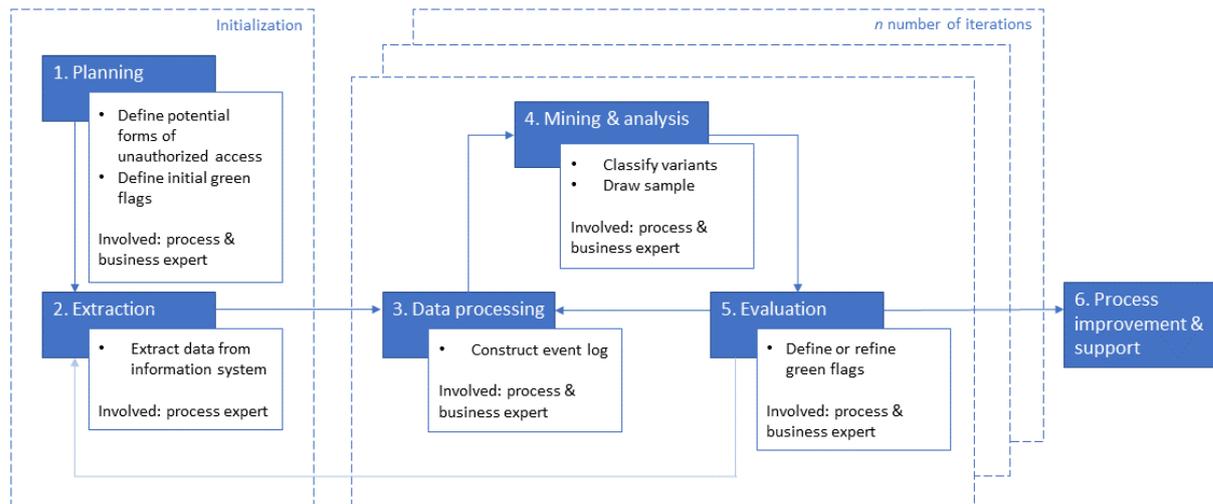


Figure 1. DEUCE: A Methodology for Detecting Unauthorized EHR Access (after PM² by van Eck et al., 2015).

1. Planning

The first objective of the planning stage is to define potential forms of unauthorized access. To determine whether some behavior may be unauthorized, initial green flags are defined. A green flag is an event that indicates a justified reason for following the path. Defining the categories of unauthorized access and an initial set of green flags is the starting point for the process mining project.

2. Extraction

In the second stage, data is extracted from the EHR system to analyze the executed process instances and separate the set into authorized and unauthorized access based on the known green flags. The extraction step may be executed multiple times in one project, as experts may require additional data to assign new green flags.

3. Data Processing

During the data processing stage, the extracted data is turned into an event log by constructing cases. Each case is an instance of a user accessing (a part of) the EHR and contains processual information around this access.

4. Mining & Analysis

In the mining & analysis stage, the event log is loaded into a process mining tool in order to explore and analyze the process. Conformance checking is carried out differently than in most studies. Instead of using a reference process model that describes the intended behavior, we look for the occurrence of green flags in the cases. When a case contains no green flags, it is classified as an instance of unauthorized access. Note that we classify cases as unauthorized, but in reality the case is *potentially* unauthorized, i.e. it needs to be further investigated by the business expert. However, for simplicity, we will refer to it as a case of unauthorized access. In each iteration, one can say with more certainty that it is in fact a case of unauthorized access.

After the application of green flags to the log, a random sample of cases is drawn and analyzed by the *business expert*. By investigating these cases in-depth and determining which of them can with certainty be confirmed as authorized and which ones cannot, the business expert provides input for the evaluation stage.

5. Evaluation

The evaluation stage concludes the analysis iteration and transitions into either the next analysis iteration or into the final stage of the project. When any cases classified as unauthorized have been determined false positives by the business expert, they will provide input for the next iteration. In collaboration with the business expert, new green flags are defined (or existing ones are refined) that

will prevent these cases and equivalent ones from being classified as unauthorized in the next iteration. This helps to reduce the number of false positives with each iteration in a structured manner.

6. Process Improvement & Support

The final stage of the methodology is reached when the process analyst and the business expert agree that the detection is mature enough. For example, this might be so when the last iteration concludes without encountering any false positives in the analysis and, thus, there is no input for a new iteration. When considered mature enough, the detection of unauthorized behavior can be implemented for continuous operational support to detect them in real-time.

4 Case Study: Detecting Unauthorized EHR Access at a Dutch Hospital

We applied the proposed methodology to a case study conducted at a large general hospital situated in the Netherlands. For the purpose of this study, we were given access to all audit logging of the EHR under the conditions of a Non-Disclosure Agreement. The hospital also provided access to documentation on their privacy policy and processes and allowed the researchers to conduct interviews with several stakeholders who were of interest to the case study.

We conducted a case study because it is a flexible research approach that is able to cope with the complex and dynamic characteristics of real world phenomena (Bryman, 2016), such as those found in healthcare (Rebuge and Ferreira, 2012). The case study uses a single-case design, as the case is common and representative (Yin, 2017). The main information system used at the case study hospital, Chipsoft HiX, is the market leader in the Netherlands; about 65% of all Dutch hospitals use it. Within the hospital, the current approach for catching unauthorized access is that the privacy officer regularly investigates a random sample of cases. However, because there are far more cases of EHR access that are authorized than unauthorized, this approach is inefficient. If a large portion of cases could automatically be determined to be authorized, the privacy officer would only need to investigate the remainder of the cases and the investigation would cover more cases of unauthorized access.

In the following sections, we describe how we applied the methodology. Although we describe the case study according to the stages of the methodology, in reality we completed four iterations of the methodology for five categories of unauthorized EHR access. During the case study, the role of process analyst was fulfilled by the first author, and the role of business expert by the privacy officer of the hospital.

4.1 Stage 1: Planning

Following the proposed methodology, we started with defining categories of unauthorized access that were to be investigated. Therefore, interviews with the Chief Information Security Officer (CISO) and the Privacy Officer (PO) were centered around how the privacy of patients is guaranteed and determining the various information security policies and practices at the hospital. The aim of the subsequent interviews with the PO, System Administrator (SA), Chief Medical Information Officer (CMIO), Chief Nursing Information Officer (CNIO), and nurse, was to discover categories of unauthorized access and how to detect those in the data. **Table 1** presents an overview of the data collection during the planning stage of the project.

Date	Stakeholder(s)	Focus of interview	Duration
5/3/2020	CISO & PO	Information security practices	1 hr
5/3/2020	PO	Introduction of EHR and work observation	4 hrs
16/3/2020	PO	Data sources, unauthorized access, and green flags	1 hr
27/3/2020	PO	Data sources, unauthorized access, and green flags	30 mins
9/4/2020	SA	Unauthorized access and green flags	30 mins

13/5/2020	Nurse	Unauthorized access and green flags	30 mins
19/5/2020	CMIO	Unauthorized access and green flags	30 mins
22/5/2020	CNIO & Nurse	Unauthorized access and green flags	1 hr

Table 1. Planning meetings with stakeholders from the hospital.

We conducted the interviews using video conference software. After receiving permission for the interview from the participant, we emailed them instructions for the interview. We sent them a single page document that explained the purpose of the study, the concept of unauthorized EHR access, and the type of information we aimed to gather from the interview with them. We kindly asked them to study the document and think about the questions for a couple of days before scheduling a call. This gave them some time to recognize unauthorized access in their work environment and come up with examples. Then, during a scheduled call, they would present their findings; we would ask follow-up questions to get a better understanding of the unauthorized EHR access that they encountered. We focused on finding out whether and how we might be able to distinguish between authorized and unauthorized access in the EHR data. The interviews were recorded and transcribed. When a type of unauthorized EHR access or green flag was mentioned, it was marked and coded as such.

4.1.1 Results

From the interviews, five categories of unauthorized EHR access were identified. **Table 2** provides an overview of the types of unauthorized access. It also shows who identified those in the interviews.

Categories of unauthorized access	PO	SA	CMIO	CNIO	Nurse
1. Unauthorized use of break the glass mechanism 'First contact'	X	X	X	X	X
2. Unauthorized use of break the glass mechanism 'ER'	X	X		X	X
3. Unauthorized use of break the glass mechanism 'Peer consultation'	X		X	X	X
4. Unauthorized access of Multimedia	X				
5. Unauthorized access of Lab result	X				

Table 2. Overview of which type of unauthorized access was mentioned by which participant.

The first three categories refer to different ways of unauthorized use of the break the glass mechanism. At the case study hospital, this mechanism is used hundreds of times a day, on average. When an employee uses the break the glass mechanism, he or she can enter a reason for doing so either by choosing one of the default options or by entering a reason in the free text field at the bottom. By checking whether the rest of the designed path was followed it can be determined whether an instance was authorized or not. For instance, consider a case for which the reason for breaking the glass is 'ER' (Emergency Room). If a patient is indeed registered at the ER shortly afterwards, then this can be seen as a main indicator that the designed path was followed in accordance with the rules. Therefore, this event can be used to set one of the green flags for this category of unauthorized access.

The two remaining categories are related to unauthorized access of EHR modules Multimedia and Lab result. These modules hold information that is sensitive by nature. These should be monitored extra carefully since they are at higher risk to be viewed by an unauthorized person. We can narrow down the investigation by finding out when it is necessary for an employee to view the module in order to carry out their job, such as when the employee is treating the patient.

To provide a better understanding of how the detection is performed in practice, we will provide a concrete example. **Table 3** provides a simplified event log of unauthorized use of break the glass mechanism with the reason 'First contact'. The event log consists of three cases that each contain one or two activities. The first case shows that shortly after the break the glass reason 'First contact' is selected, a green flag is detected, namely *Appointment with patient made*. Because the case contains a green flag, it is determined authorized, as shown in the Class column. In case 2, we have no way to justify this event, thus it is labeled as unauthorized. The last case shows that some green flags, in this

case *Referral of patient registered*, can also occur before the break the glass mechanism is used, as long as they fall within the specified time window.

In order to classify cases, we take different process perspectives into account (Mannhardt et al., 2016). First of all, the perspective most often used in process mining: the *control flow* perspective. Green flags in this category relate to the order of activities found in the data. For instance, in the example in Table 3, we look for cases where first the break the glass mechanism is used, and then an appointment is made. A second perspective that we take into account for the example in Table 3 is the *time* perspective. Case 1 is only classified when the time between the appointment and the use of the break the glass mechanism is below a certain threshold. Third, a number of green flags are related to the *resource* perspective. This is most evident in the example from the introduction, where we take into account which doctor accesses a patient’s record and whether or not he or she also treating the patient. Last, some green flags focused on the *data* perspective. This includes situations where we base the classification on additional data attributes.

Case ID	Timestamp	Activity	Resource	Class
1	10:04	Break the glass mechanism ‘First contact’ used	A	Authorized
	10:07	Appointment with patient made	B	
2	11:42	Break the glass mechanism ‘First contact’ used	C	Unauthorized
3	12:35	Referral of patient registered	D	Authorized
	15:49	Break the glass mechanism ‘First contact’ used	D	

Table 3. Simplified example of event log for unauthorized EHR access category 1.

4.2 Stage 2: Extraction

The second stage is aimed at extracting the data necessary to analyze the instances of unauthorized EHR access. The case study hospital's EHR (Chipsoft HiX) is a Process-Aware Information System, which provides very detailed logs about the activities that have been executed in the system (Mans et al., 2008; van der Aalst, 2016). The event data needed to be manually extracted from the EHR system, which was done in accordance with the principle of data minimization so that only data was collected that was necessary for detecting unauthorized EHR access. Due to the iterative improvement of the detection, this meant that data extraction needed to be repeated several times. The data was extracted from the database using SQL queries and analyzed on a secure dedicated server for the duration of the case study. Only pseudonymized data was extracted from the system.

We opted for January and February of 2020 as the scope of the data for several reasons. The data is recent, yet it is minimally influenced by the COVID-19 outbreak in the Netherlands, which commenced in March 2020. This slice of the data contains thousands of instances for each category of unauthorized EHR access, providing a sufficiently large dataset to create a proof of concept of the detection.

4.3 Stage 3: Data Processing

After extracting the data, we pre-processed it by transforming the data into events, filtering out test patients, relating sub-specialisms to main specialisms, and relating patients to their aliases in case of duplicate registrations. Then, we wrote custom scripts using the R programming language for statistical computing (R Core Team, 2013) to assign the events to cases in order to construct event logs. This assignment was performed based on characteristics such as the timestamp, patient number, employee number, employee function, and the employee’s specialism. The conditions for these characteristics varied for each type of event and the situation of the patient at that moment.

4.4 Stage 4: Mining & Analysis

To analyze the event logs, we used the process mining tool Disco by Fluxicon (<https://fluxicon.com/disco>) since it is lightweight and intuitive to use yet contains all the required functionalities. Based on the known green flags in each iteration, we distinguished the authorized from the unauthorized EHR cases. At the start, little green flags are known. Thus, a large portion of the cases is assumed unauthorized. Consequently, the number of false positives starts off relatively high, and decreases with each iteration.

In each iteration, we drew ten random cases. The PO would then determine whether these cases were false positives, and if yes, add new green flags to the list. For each category of unauthorized EHR access, four iterations were completed in this way during twelve meetings with the PO between April and August 2020 (**Table 4**). The meetings took 66 minutes on average.

Note that the iterations of the five categories are not aligned in time. For example, categories 4 and 5 were only discovered after the second iterations of categories 1 and 2 had finished. As green flags identified in one category may be reused for analyzing cases in another category, it is important to take this into account when reflecting on the scores presented in the following section.

Date	3/4	24/4	7/5	14/5	27/5	5/6	17/6	26/6	10/7	17/7	27/7	5/8
CAT1	1			2		3				4		
CAT2		1		2			3		4			
CAT3			1			2			3			4
CAT4					1			2			3	4
CAT5					1			2			3	4

Table 4. Overview of iterative cycles (dates all in 2020).

4.5 Stage 5: Evaluation

The aim of the evaluation stage of the methodology is to determine whether to move to the Process Improvement & Support stage or to continue with another iteration. In order to determine the accuracy of the detection, we consider two metrics:

1. Precision

Precision and recall are metrics that are often used to determine the accuracy of a classification. Precision can be calculated based on the results from the sample of cases investigated in stage 4. We define precision as follows:

$$\text{Precision score} = \text{true positives} / (\text{true positives} + \text{false positives})$$

We do not include recall, as investigating true and false negatives requires much more effort than true and false positives. We discuss this more in detail in the discussion.

2. Percentage of unauthorized access

We also collect the total number of assumed unauthorized cases for each category. By relating this number to the total number of cases we get the percentage of unauthorized access:

$$\text{Percentage of unauthorized access} = \text{unauthorized cases} / \text{total cases} \times 100\%$$

When the percentage no longer changes substantially over the iterations, this tells us something about the maturity of the detection. During the analysis stage, new green flags may still be added, but they do not have a major effect on the classification of cases. In other words, the curve flattens, which means that the accuracy of the classification cannot be much further improved.

4.5.1 Results

Table 5 provides the full results on the five categories of unauthorized access and corresponding iterations. Below the table, we discuss the results using the two metrics.

General			Cases				Sample		
Category	Iteration	# Green flags	% Unauthorized	Unauthorized cases	Authorized cases	Total number of cases	True positives	False positives	Precision score
1	1	2	43.8	9519	12213	21732	3	7	0.3
	2	6	30.3	6585	15147		7	3	0.7
	3	7	22.7	4933	16977		9	1	0.9
	4	7	20.3	4412	17320		8	2	0.8
2	1	1	43	2095	2778	4873	4	6	0.4
	2	3	26.6	1296	3577		5	5	0.5
	3	6	17.3	843	4030		6	4	0.6
	4	8	15.8	770	4103		7	3	0.7
3	1	3	58.3	2123	1519	3642	7	3	0.7
	2	8	49.9	1817	1825		7	3	0.7
	3	9	47	1712	1930		9	1	0.9
	4	12	42.6	1551	2091		9	1	0.9
4	1	2	50.6	4540	4432	8972	4	6	0.4
	2	8	34	3050	5922		4	6	0.4
	3	14	27.5	2467	6505		8	2	0.8
	4	15	25.5	2288	6684		10	0	1
5	1	1	48	11450	12405	23855	0	10	0
	2	7	20.1	4795	19060		5	5	0.5
	3	12	14.3	3411	20444		7	3	0.7
	4	14	13.1	3125	20730		8	2	0.8

Table 5. Results per category.

Figure 2 shows the precision score and percentage of unauthorized access for each category across the iterations. The precision score is calculated on the basis of the business expert’s investigation of a sample of cases. It relates to the cases currently classified as unauthorized, and the extent to which these are correctly classified as such. As the curve approaches the maximum precision score of 1.0, there are very little false positives left which also means little input for improvement. This is an indicator that the detection has nearly been optimized.

The percentage of cases classified as unauthorized decreases with each iteration, as can be expected because more and more green flags are applied to the data. As the detection of unauthorized access matures, the curve flattens, meaning that most of the false positives have been eliminated as far as possible. Note that some categories of unauthorized access might occur more often than others, meaning that their percentage will be higher even when considered mature.

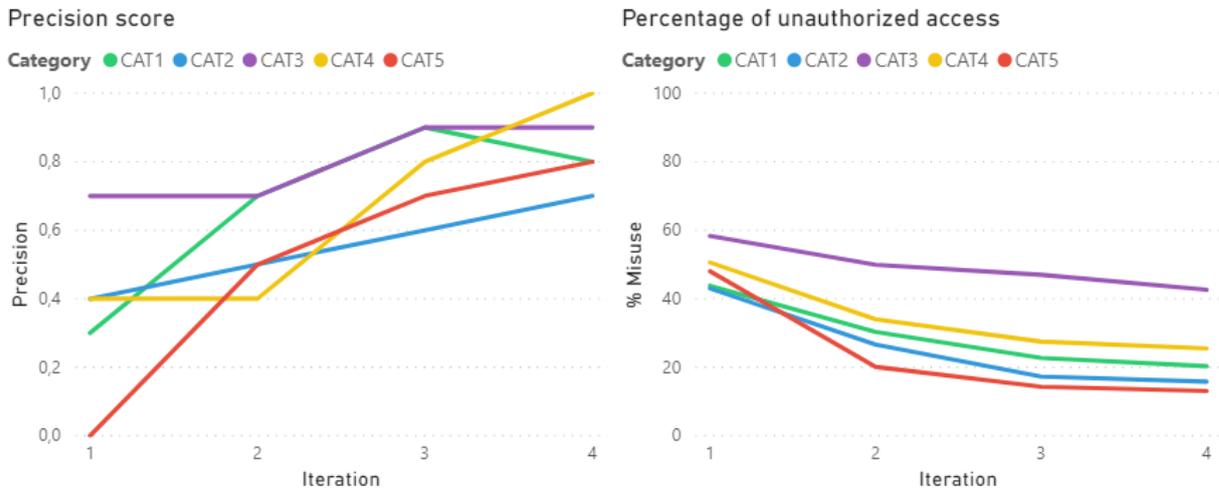


Figure 2. Precision score and percentage of unauthorized access.

The combination of metrics gives a more complete picture of the maturity of the detection as opposed to only looking at either one of them. For example, from the precision score figure, it looks as though the detection of category 3 (the purple line) did not improve between the first and second iteration. However, this may be skewed by the sample of the first iteration having an over-representation of true positives. Ground for this suspicion is that 0.7 is an unusually high precision score for the first iteration, when comparing it to the other categories. In the unauthorized access percentage figure it can be observed that the percentage decreased by nearly 10%. In other words, many false positives were eliminated so there was in fact a large improvement in the second iteration.

4.6 Stage 6: Process Improvement & Support

After four iterations for each category of unauthorized access, the detection was considered mature enough for the purpose of this project. As a result of this study, the techniques used and green flags developed have been implemented in the monitoring software used by the hospital. As such, the Privacy Officer no longer has to draw samples from the complete set of cases, but will be focusing on those cases that do not contain green flags.

5 Discussion

In this study, we proposed DEUCE: a methodology aimed at identifying authorized behavior such that unauthorized behavior remains and can be further investigated. We evaluated the methodology empirically in the form of a case study, using real data from a hospital. In the next sections, we discuss how this approach differs from the traditional approach used by the hospital, how it relates to the approaches currently used in conformance checking, and how its principles can be applied outside the detection of unauthorized EHR access.

5.1 DEUCE versus the Hospital's Traditional Approach

Before undertaking this study, the hospital's way of detecting unauthorized EHR access was by manually checking a sample of EHR views by employees. A PO would try to contextualize the cases one by one: checking what occurred before and after the suspicious view, who the user was, what role he or she has within the organization, etc. However, the effort needed to investigate all these cases is high. For each sample of cases that is checked, the effort required remains equally high. In contrast, our proposed methodology introduces an automated component. Green flags are defined by the privacy officer and used to automatically classify all cases in the set as either authorized or unauthorized. Consequently, the list of green flags is appended with every new detection. Therefore, the effort needed for the next detection decreases, rather than remaining the same.

In recent years, several techniques for conformance checking have been proposed, many of which relate to rule checking, token replay, or alignments (Carmona et al., 2018). Before these techniques can be applied, an event log and process model need to be present. In some lines of business, process models are readily available. However, in healthcare, this is not often the case. Process models in healthcare often provide an idealized version of the process, while deviations from these models are frequent and many times accepted (Lu et al., 2014). Applying conformance checking techniques on event logs with these models would result in a large portion of behavior classified as deviations, providing a false reflection on reality. Our contribution to the field is that we applied conformance checking techniques and concepts to processes that are loosely structured. We set off from the idea that in some situations, it is easier to say with certainty that some case conforms to the model, while it is much more difficult to say with certainty that some case deviates from the model. This may not merely apply to the detection of unauthorized EHR access, but can be applied in a much broader sense to processes that are loosely structured and difficult to model.

5.2 Broader Applications

Although our methodology is described in relation to the detection of unauthorized EHR access and applied to a case study in healthcare, we believe the principles can be applied in a broader sense. First, it can be applied to enable conformance checking of loosely structured processes in general, such as in relation to the clinical pathways mentioned in the previous section (Carmona et al., 2018). However, the need for flexibility is not constrained to the healthcare sector. There are many other processes that benefit from a flexible setup. For example, knowledge-intensive business processes in general are difficult to fully anticipate in advance (Unger et al., 2015). When applying conformance checking to this type of processes, one may benefit from taking the principles from our methodology into account; drawing samples of cases to define green flags, applying those green flags in order to arrive at unusual cases, and identifying opportunities to improve the process.

Second, our approach to discover suspicious behavior could also be helpful in the context of detecting workaround behavior, i.e. users deviating from the intended way of using information systems (Ejnefjäll and Ågerfalk, 2019). Recently, some studies have attempted to use process mining to detect workarounds (Outmazgin and Soffer, 2016; Beerepoot et al., 2020; Weinzierl et al., 2020). These studies also build on specifications of patterns of behavior that indicate the occurrence of workarounds. Arriving at these specifications requires knowledge from business experts, collected during iterative meetings (Beerepoot et al., 2020). Therefore, the steps outlined in our methodology may also be of use in distinguishing workaround behavior from behavior intended by the system's designers and managers.

Last, we believe our approach can be of help in applying conformance checking to processes that change over time, a field of study known as *concept drift*. Within process mining research, concept drift studies are concerned with how processes change over time and whether these changes can be detected and explained (Bose et al., 2011). As processes change, models documenting those processes cannot be expected to always change with them, if there are models at all. When studying processes that have changed over time, the principles of our methodology may also apply. For example, one might draw a sample from each time period, manually check whether the behavior in the log conforms with the expected behavior in that time frame, classify this behavior as either conforming or non-conforming, and apply the green flags to all behavior in that time period.

5.3 Limitations

In this study, we have been fortunate to be able to systematically study five categories of unauthorized EHR access with a business expert, discussing ten cases for each category over four iterations. Doing so over several iterations has contributed towards improving the reliability of the score. However, the approach needs to be applied on a larger scale to more precisely evaluate the improvement of results over time. Moreover, we have focused on the reduction of false positives and not false negatives. The main reason why our methodology puts the emphasis on false positives is that they are easier to detect

than false negatives. False negatives remain hidden in the larger portion of the cases that has already been systematically determined as authorized due to one or multiple green flags. When analyzing a sample of assumed unauthorized cases, one is almost certain to encounter false positives, especially during the earlier iterations. Whereas when analyzing a sample of authorized cases, it is less likely that false negatives will be encountered. The reason for this is that there are far more true negatives than true positives. In future work, analyzing a random sample of cases from the set of green flags, could be included in the methodology to determine the number of false negatives. This would also allow for calculating the fitness of the detection. However, detecting false negatives does require more effort from the business expert. When time of business experts is limited, which it often is, it may be wiser to invest in the more efficient reduction of false positives.

6 Conclusion

As a result of the loosely structured and ad-hoc nature of healthcare processes, distinguishing the ‘right’ way of doing things from the ‘wrong’ way, is complex, especially when it comes to detecting EHR access. The currently popular approach, which comes down to manually checking access, is labor-intensive and limited in the sense of cases that can be checked. In this paper, we studied how process mining can be used to support the detection of unauthorized EHR access and propose the DEUCE methodology. To arrive at the methodology, we drew on the principles of conformance checking, the red flag approach, and the principle of exclusion. Moreover, we adapted the stages from the well-established Process Mining Project Methodology, including the iterative analysis cycles and the involvement of business experts throughout the project. We applied the methodology in a case study at a Dutch hospital, systematically classifying five types of access behavior into authorized and unauthorized cases. We showed how, over multiple iterations and for each category, the precision of the detection increased and the classification of unauthorized access substantially matured. Our methodology provides a direct improvement with respect to the traditional approach of the hospital, with regards to both the effort necessary to detect unauthorized cases and the completeness of the detection in terms of the cases checked. As a consequence of our endeavor, the hospital involved has implemented the developed principles and classifications and is now using the approach to more efficiently catch unauthorized EHR access and to better ensure the privacy of patients.

References

- Albrecht, W. S., C. O. Albrecht, C. C. Albrecht and M. F. Zimelman. (2011). *Fraud examination*. Cengage Learning.
- Baader, G. and H. Krcmar. (2018). “Reducing false positives in fraud detection: Combining the red flag approach with process mining.” *International Journal of Accounting Information Systems*, 31, 1–16.
- Beerepoot, I., X. Lu, I. van de Weerd and H. A. Reijers. (2020). “Seeing the signs of workarounds: a mixed-methods approach to the detection of nurses’s process deviations.” In: *Hawaii International Conference on System Sciences*.
- Borycki, E., R. S. Joe, B. Armstrong, P. Bellwood and R. Campbell. (2011). “Educating health professionals about the electronic health record (EHR): Removing the barriers to adoption.”
- Bose, R. P. J. C., W. M. P. Van Der Aalst, I. Žliobaite and M. Pechenizkiy. (2011). “Handling concept drift in process mining.” In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Vol. 6741 LNCS, pp. 391–405). Springer, Berlin, Heidelberg.
- Bryman, A. (2016). *Social research methods*. Oxford university press.
- Carmona, J., B. van Dongen, A. Solti and M. Weidlich. (2018). *Conformance Checking*. Springer.
- Ejnefjäll, T. and P. J. Ågerfalk. (2019). “Conceptualizing Workarounds: Meanings and Manifestations in Information Systems Research.” *Communications of the Association for Information Systems*, 45(1), 20.

- Homayounfar, P. (2012). "Process mining challenges in hospital information systems." In: *2012 Federated Conference on Computer Science and Information Systems (FedCSIS)* (pp. 1135–1140).
- Koppel, R., S. Smith, J. Blythe and V. Kothari. (2015). "Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient?" In: *Studies in Health Technology and Informatics* (Vol. 208, pp. 215–220).
- Lovis, C., S. Spahni, N. Cassoni and A. Geissbuhler. (2007). "Comprehensive management of the access to the electronic patient record: Towards trans-institutional networks." *International Journal of Medical Informatics*, 76(5–6), 466–470.
- Lu, X., R. S. Mans, D. Fahland and W. M. P. van der Aalst. (2014). "Conformance checking in healthcare based on partially ordered event data." In: *Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA)* (pp. 1–8).
- Mannhardt, F., M. De Leoni, H. A. Reijers and W. M. P. Van Der Aalst. (2016). "Balanced multi-perspective checking of process conformance." *Computing*, 98(4), 407–437.
- Mans, R. S., W. M. P. Van der Aalst and R. J. B. Vanwersch. (2015). *Process mining in healthcare: evaluating and exploiting operational healthcare processes*. Springer.
- Murphy, A. R., M. C. Reddy and H. Xu. (2014). "Privacy practices in collaborative environments: a study of emergency department staff." In: *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (pp. 269–282).
- Outmazgin, N. and P. Soffer. (2016). "A process mining-based analysis of business process workarounds." *Software & Systems Modeling*, 15(2), 309–323.
- Ramezani, E., D. Fahland and W. M. P. van der Aalst. (2012). "Where did I misbehave? Diagnostic information in compliance checking." In: *International conference on business process management* (pp. 262–278).
- Rebuge, A. and D. R. Ferreira. (2012). "Business process analysis in healthcare environments: A methodology based on process mining." *Information Systems*, 37(2), 99–116.
- Rodrigues, J. J. P. C., I. de la Torre, G. Fernández and M. López-Coronado. (2013). "Analysis of the security and privacy requirements of cloud-based electronic health records systems." *Journal of Medical Internet Research*, 15(8), e186.
- Rojas, E., J. Munoz-Gama, M. Sepúlveda and D. Capurro. (2016). "Process mining in healthcare: A literature review." *Journal of Biomedical Informatics*, 61, 224–236.
- Rojas, E., M. Sepúlveda, J. Munoz-Gama, D. Capurro, V. Traver and C. Fernandez-Llatas. (2017). "Question-driven methodology for analyzing emergency room processes using process mining." *Applied Sciences*, 7(3), 302.
- Unger, M., H. Leopold and J. Mendling. (2015). "How much flexibility is good for knowledge intensive business processes: A study of the effects of informal work practices." In: *System Sciences (HICSS), 2015 48th Hawaii International Conference on* (pp. 4990–4999).
- van der Aalst, W. (2016). *Process Mining: Data Science in Action*. Springer.
- van Eck, M. L., X. Lu, S. J. J. Leemans and W. M. P. van der Aalst. (2015). "PM²: A Process Mining Project Methodology." In: *International Conference on Advanced Information Systems Engineering* (pp. 297–313).
- Weinzierl, S., V. Wolf, T. Pauli, D. Beverungen and M. Matzner. (2020). "Detecting Workarounds in Business Processes-a Deep Learning method for Analyzing Event Logs." In: *ECIS*.
- Yin, R. K. (2017). *Case study research and applications: Design and methods*. Sage publications.
- Zerbino, P., D. Aloini, R. Dulmin and V. Mininno. (2018). "Process-mining-enabled audit of information systems: Methodology and an application." *Expert Systems with Applications*, 110, 80–92.