Towards a Standard Process enabling AI-support for Safety and Conformity of Medical Devices

Peter Pfeiffer[1,2], Heike Sander[3], Peter Fettke[1,2] and Wolfgang Reisig[4]

[1]German Research Center for Artificial Intelligence (DFKI), Campus D3.2, 66123 Saarbrücken, Germany
[2]Saarland University, 66123 Saarbrücken, Germany
[3]DHC Business Solutions, Landwehrplatz 6-7, 66111 Saarbrücken, Germany
[4]Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany

Abstract

Within the lifecycle of a medical device, Post-Market Surveillance (PMS) refers to the proactive, systematic and periodic monitoring of adverse events and other findings related to medical devices on the market, aiming to ensure the safety of patients, healthcare professionals and other users. PMS and associated activities, like Vigilance, are required by law and part of the obligatory risk management for medical products where regulatory requirements to be met are, for instance, defined by the Medical Device Regulation. Information to be monitored includes adverse event reports, clinical studies, scientific publications and other reports from any relevant internal or external data source. This requires the collection of heterogeneous data from various data sources in order to perform qualitative and quantitative analysis. In this work, we propose an initial version of a standard process with emphasis on data collection and analysis of data relevant for PMS and associated activities within the medical device lifecycle which is based on experience from regulatory specialists in this field. In addition, the challenges and potential for AI-support within these processes are discussed.

Keywords

Post-Market Surveillance, Vigilance, Business Process Management

1. Introduction

Post-Market Surveillance (PMS) is statutory for European medical device manufacturers in order to ensure the safety of patients, users and third parties [1]. PMS contains a proactive, legally obligatory, regularly and continuous process of ensuring the safety of medical devices after their launch on the market. One of the central aspects is the collection and evaluation of data on product safety with the aim to detect possible issues with a medical device as soon as possible. The requirement of defining and maintaining PMS processes is part of the Medical Device Regulation (MDR) [1], a European standard framework for medical device safety introduced in May 2021. The introduction of MDR increased the requirements for PMS both quantitatively and qualitatively; new processes to meet those requirements are still emerging. PMS includes the active observation of the device itself, similar devices, and devices with similar parts within the lifecycle of the medical device. It also includes the active monitoring and evaluation of

The International Health Data Workshop (HEDA-2022) in conjuction with Petri Nets 2022.
peter.pfeiffer@dfki.de (P. Pfeiffer); heike.sander@dhc-vision.com (H. Sander); peter.fettke@dfki.de (P. Fettke); reisig@informatik.hu-berlin.de (W. Reisig)

clinical studies and notifications about the devices, adverse events, and complaints from any relevant source as well as trend analysis [2]. The MDR summarises this requirement as the inclusion of "any relevant data" [1]. Within the lifecycle of a medical device, PMS is repeated annually or biannually. Vigilance [2], one of the associated processes for medical device safety which is also required and described in the MDR, shares, among others, the data collection and analysis with PMS. While the Post-Market Surveillance is a proactive system and repeated periodically, Vigilance is an ongoing reactive system. Vigilance reacts to all reported and experienced incidents, both from PMS and other sources, and includes a trend reporting as well [2]. They are intended to improve the protection of patients, healthcare professionals and other users of medical devices by reducing the likelihood of re-occurrence of incidents related to the use of a medical device.

Post-Market Surveillance, Vigilance and other associated processes mandatory by MDR have in common that dispersed data from various sources has to be collected, reviewed and analysed. Searching for relevant data has to follow a standardized procedure and all activities must be documented. Using the relevant data, different reports are created and sent to notified bodies[1]. As the MDR mandates the inclusion of any relevant data, there are several main data sources of interest. Data is unstructured multilingual texts; they come from unsolicited sources (e.g., spontaneous reports from healthcare professions and patients; screenings of global and local literature; screenings of internet/digital media, internal complaints, etc.) or solicited sources (e.g., derived from organised data collection system such as clinical trials, non-interventional studies, registries). For clinical studies, Pubmed[2] and Medline[3] are predominantly used. Adverse events are mainly included within country specific databases, for instance FDA Maude[4] in the USA, BfArM[5] in Germany and Swissmed[6] in Switzerland. Almost every country has its own database, typically in the country's language. For Europe, the deployment of the EUDAMED[7] database is planned, which is meant to collect all adverse events occurring in all the European countries and provide them in all European languages.

While there are many regulations and guidelines to follow, to the best of our knowledge, there are no standard processes in the form of a process model for such activities. For instance, there is no standard procedure that supports regulatory specialists in obtaining insights that can be used for writing a conformity report. From our experience, this makes it hard for manufacturers to follow the MDR and conform to all requirements imposed thereby. Furthermore, PMS and Vigilance are time-consuming tasks which are often not supported by information systems. Finding relevant documents and reports is a challenge due to the distribution of the data in different databases in different languages, different data formats and the huge amount of reports that are potentially of interest. Thus, there is a lot of manual work to be done in order to collect, review and analyse relevant information.

In this paper, we propose a scheme for performing such processes in a regulatory environment.

[1] In the EU they are designated in each country by the member states[1] and differ from country to country
[2] https://pubmed.ncbi.nlm.nih.gov/
[3] https://www.nlm.nih.gov/
[4] https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfMAUDE/search.CFM
[5] https://www.bfarm.de/
[6] https://www.swissmedic.ch
[7] https://ec.europa.eu/tools/eudamed/

It serves as a basis for processes containing data collection and analysis of medical device data to support companies fulfilling the obligations imposed by regulations. For instance, to help regulatory specialists to collect and analyse data in a systematic way for writing reports that have to be submitted to notified bodies. The proposed scheme is split in several modules that can be re-used and specified for different needs. For instance, the module for collecting and reviewing data sources can be used in various processes that require the analysis of medical device data. As the scheme also defines what data is used, processed and created, it can be used as a blueprint for building software that supports these processes. Furthermore, some time consuming and repetitive tasks can be automated or supported by means of artificial intelligence (AI) to support the user in finding all relevant information. Thus, we propose challenges and time consuming tasks that regulatory specialists face when performing such processes and demonstrate how AI can be integrated to make the data collection and analysis more efficient and effective, which is supposed to support the safety of medical devices. At the same time, the processes should comply with regulatory requirements. As different companies perform these processes differently and regulatory requirements are subject to change in the future, we aim for a "best practice" scheme that serves as an initial version for future work.

The remainder of this paper is structured as follows. In the next Section, we will introduce the scheme for PMS and the associated processes. In Section 3, we will present challenges and possibilities to give AI-support within the processes. Section 4 closes the paper with a discussion and conclusion.

2. A Scheme supporting the Safety and Conformity

How PMS and similar processes are performed varies from company to company, and product to product. However, we observed similar sub-procedures how users proceed. In this section, we will first introduce our methodology before introducing the runs we have observed in practice and the scheme that is derived from the runs.

2.1. Methodology and Techniques

We started by collecting regulatory requirements for PMS and similar processes such as laws, regulations and standards, e.g., proposed by MDR. Next, we conducted several interviews with regulatory specialists of three international manufacturers of medical devices that are based in Germany to find out how they proceed, e.g., in PMS. The companies differ in size and have devices of different classes on the market. From the interviews we obtained several sample runs [3] showing different ways how to perform safety and conformity processes in practice as well as further insights into regulatory requirements. Using the runs, we derived and modelled detailed processes which are split into higher-level and lower-level modules. While the higher-level modules structure the lower-level ones, the lower-level modules contain detailed information about which activities to perform as well as what data and information is used, processed and created. This gives a structured and hierarchical view on the relevant activities and allows a flexible adjustment and adaptation to specific needs of different stakeholders. For instance, the inner of each module can be changed without destroying the interoperability with other modules.

We use HERAKLIT [4] for modelling, which provides an integrated view on structure, data and local behaviour of large scale system. HERAKLIT makes extensive use of Petri nets and extends it with concepts such as modularization, local and global behaviour as well as algebra for describing the data[8]. Each run [3] presented in this section (Fig. 1-3) shows one example how such processes are performed - rectangles represent actions that have happened while circles represent data that has been created or consumed by actions. On the other side, modules describe systems on a more abstract level. Arrows indicate what data is send between modules. The inner of modules consists of Petri nets with data carrying tokens that show the detailed behaviour of the module. The signature indicates the structure of the data.

Within the interviews with the regulatory specialists we collected and summarized challenges they face. To make their work easier and more efficient, we added possibilities how to support them with methods based on artificial intelligence.

Afterwards, we presented the runs and the scheme to the regulatory specialists and interviewed them again to evaluate its correctness and usefulness. The specialists agreed that the runs are plausible and the scheme covers the processes they perform. Furthermore, they found no other problems with the models. We concluded that the scheme has no obvious errors and fits the procedure of how PMS and similar processes are performed in practice. Although the scheme has no errors according to the specialists, this does not guarantee that a report created when following the scheme will be approved by the notified bodies in the end, as the approval depends on the reports content.

2.2. Medical Device Lifecycle

In the lifecycle of a medical device there is a pre-market and a post-market phase [2]. While the pre-market phase consists of activities required before bringing a medical device on the market, the post-market phase monitors and ensures the safety of a medical device on the market. To introduce a device on the market, an initial conformity assessment takes place, which also contains a review of the technical documentation from the manufacturer on the safety and performance of the device. We focus on the medical device types that require the involvement of a notified body (classes I*, IIa, IIb, III), yielding in a CE mark. Conformity assessment refers to the process of demonstrating whether the requirements of the MDR relating to a device have been fulfilled [1]. Among others, this contains a clinical evaluation of the device and a PMS for similar devices. The clinical evaluation includes clinical studies on the device as well as relevant reports from similar devices or devices with similar parts. In contrary, the PMS is intended to discuss the behavior of similar devices on the market and assess the associated risks. This involves researching similar devices that are already on the market. The collected data is used for the initial conformity assessment, which yields a conformity report [2].

Once the device is on the market, the data collection continues. The requirements are described in detail in Annex III of the MDR [1]. Points to address within PMS are for instance detection of adverse events, comparison of a new device's performance against current standards of care, compliance with regulatory requirements of a medical device, and continuous monitoring of safety and effectiveness of the device in the intended patient population. The details of how to

[8]We refer the reader to [4] for a detailed explanation of the concepts used in HERAKLIT

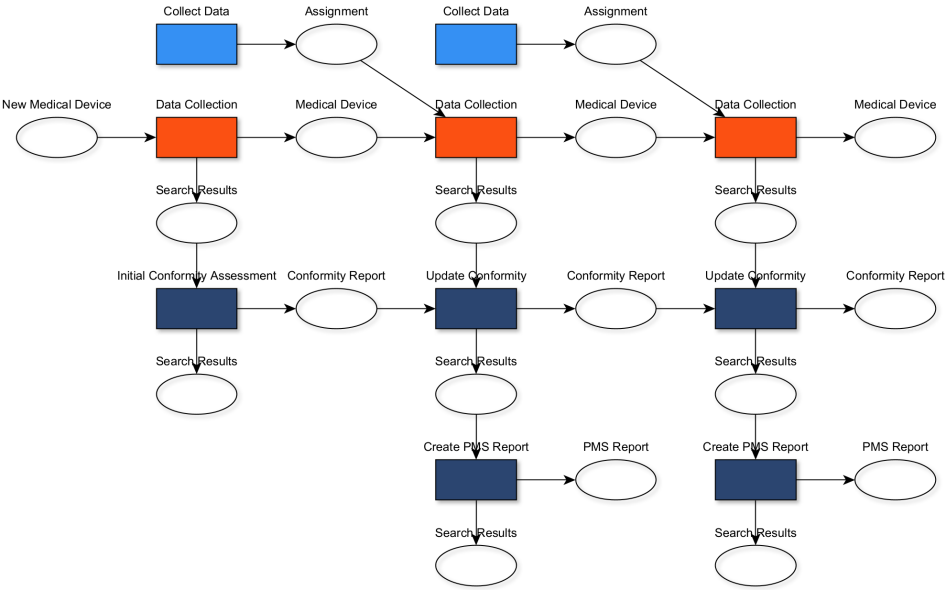

Figure 1: A high-level exemplary run of a medical device's lifecycle containing three data collections and additional actions.

perform PMS have to be described in the PMS Plan, which is part of the technical documentation a company has to implement for each medical device. The plan outlines the strategy to collect and analyze data, the follow-up strategy on collected complaints, the communication of information to regulators and users, and the corrective actions on devices, if necessary. It contains also the strategy to produce a Post Market Clinical Follow-Up (PMCF) plan, or a rationale why it is not required. A PMCF confirms the clinical performance and safety of the device and ensures continued acceptability of benefit/risk ratio of the device. The PMS yields either a PMS report (Device Class I) or a Periodic Safety Update Report (PSUR) (Classes IIa, IIb, III), depending on the classification of the device according to its involved risk.

In the following, we will show three sample runs that demonstrate what actions are performed within the lifecycle of a medical device and what data objects are created. Note that these are just examples of what could have happened. Many other ways are possible and plausible. Different colors indicate which module a action is part of.

2.2.1. Run with no Intervention

Figure 1 shows an exemplary run of the actions performed in the lifecycle of a medical device i.e., what activities have been performed and what objects are created. First, *Data Collection* is performed to obtain the information, i.e., the search results used for writing the initial conformity report. With the conformity report, the device can be released to the market. Eventually, an assignment indicates to perform the *Data Collection* again, which is triggered by action *Collect Data*. The search results are used to update the conformity report and create a PMS report. At this point the collected information can also be used to create other reports required by MDR like a trending report or PSUR. The same procedure of *Data Collection* and associated actions is

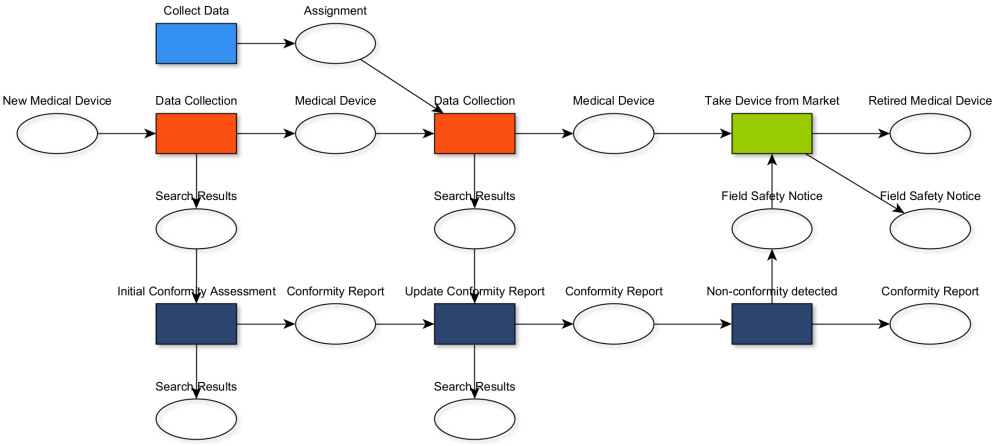

Figure 2: A high-level exemplary run of a medical device's lifecycle with retiring the device.

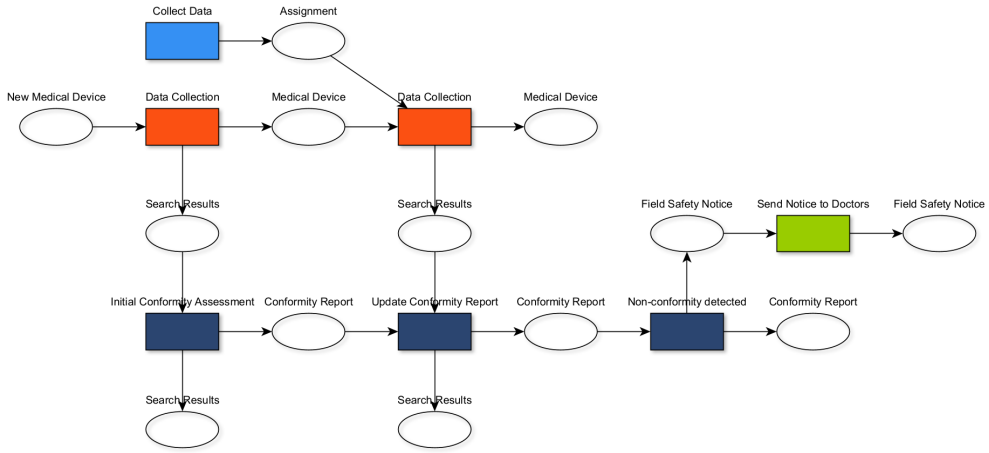

Figure 3: A high-level exemplary run of a medical device's lifecycle with sending a notification.

done again after a certain time. This particular run ends after two iterations.

2.2.2. Run with Non-conformity and Retirement

In contrast to the previous one, the run in Figure 2 shows the situation where a non-conformity was detected while the device is on the market. In such an event, the device may be retired which requires an update of the initial conformity report with information from the second data collection. Afterwards, a notification is sent and the device is retired from the market.

2.2.3. Run with Field Safety Notice

Figure 3 shows the situation where insights gathered during data collection have to be communicated to notified bodies and users. This may be required if a problem has been detected that

is not as serious as to retire the device but requires to send a Field Safety Notice to all users to prevent the problem in the future.

2.3. Modularization

The previous three runs have shown some of the behaviour that can happen in the lifecycle of a medical device. The runs can be structured into HERAKLIT modules to capsule their behaviour and create a level of abstraction which follows the idea of reference modelling [5]. Figure 4 shows the module *Medical Device Lifecycle* which consists of five sub-modules that structure the activities within the lifecycle of a medical device. The five submodules *Assignments*, *Data Collection*, *Additional Actions*, *Devices in Market* and *Devices out of Market* contain modules that support the activities in the lifecycle of a medical device. Different activities are performed in modules and data is exchanged between modules. For simplicity, the inner of the modules is yet not shown. However, each action in the runs in the previous figures is part of a certain module, indicated by the color-coding. For instance, the action *Collect Data* in Figure 1 is part of the module *Assignments*.

The interactions are as follows. If, for instance, a PMS is due, module *Assignments* sends an *Assignment* to module *Data Collection* which starts the procedure to collect relevant data. *Devices in Market* contains all devices in market and those that are to be put on the market. All relevant data obtained in *Data Collection* is passed to module *Additional Actions* as *Search Results* where additional analysis are performed, and reports, such as PMS-report or action plans, are created. Notices are sent back to *Devices in Market* to trigger actions like sending a Field Safety Notice to the notified bodies, removing the device from the market or notifying users of the medical device. If a device has been retired, it is moved to module *Devices out of Market*.

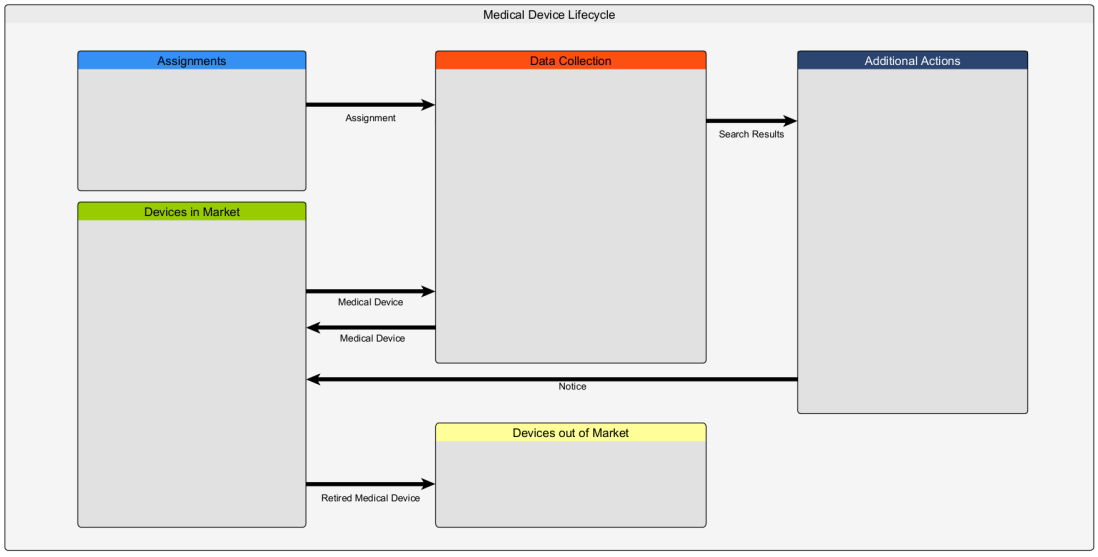

Figure 4: The lifecycle of a medical device modelled as HERAKLIT scheme.

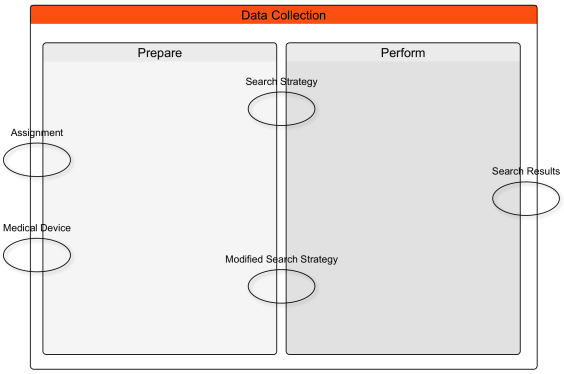

Figure 5: Module *Data Collection* and its two sub-modules.

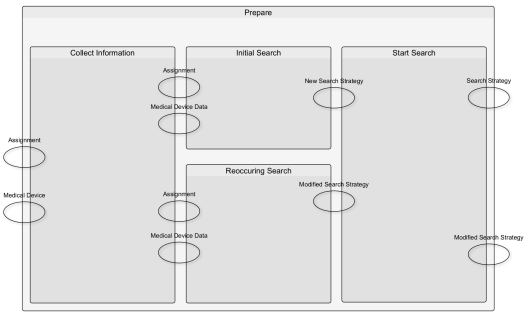

(a) The module *Prepare*

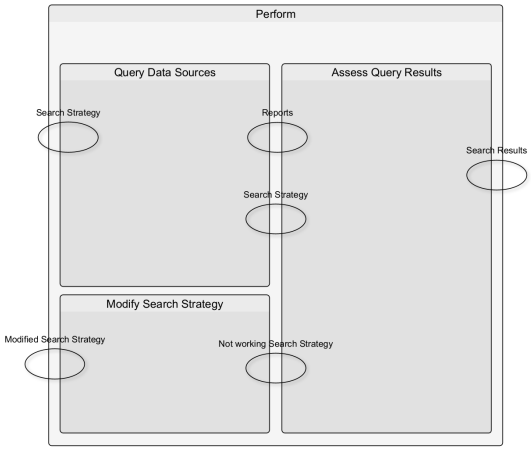

(b) The module *Perform*

Figure 6: The two submodules of module *Data Collection* in detail.

2.4. Data Collection and Additional Actions

The actions to actually collect data, assess it for relevance and create reports are performed in the modules *Data Collection* and *Additional Actions* which we will focus on now. Figure 5 shows the inner of the module *Data Collection* which is again divided into two modules shown in Figure 6. While in *Prepare*, the search strategy is defined, the actual search and assessment of relevant reports is performed in module *Perform*. Both modules contain sub-modules again. In submodule *Prepare*, a decision is made whether the assignment requires to perform an initial data collection for a new product or for a known device. If it is an initial data collection, a new search profile will be created in module *Initial Search*. If a search strategy exists for this medical device, module *Reoccurring Search* is used to find the corresponding search profile and update it (if required). Module *Start Search* combines both situations as well as the situation where the search strategy is passed back from an unsuccessful search (indicated by *Modified Search*

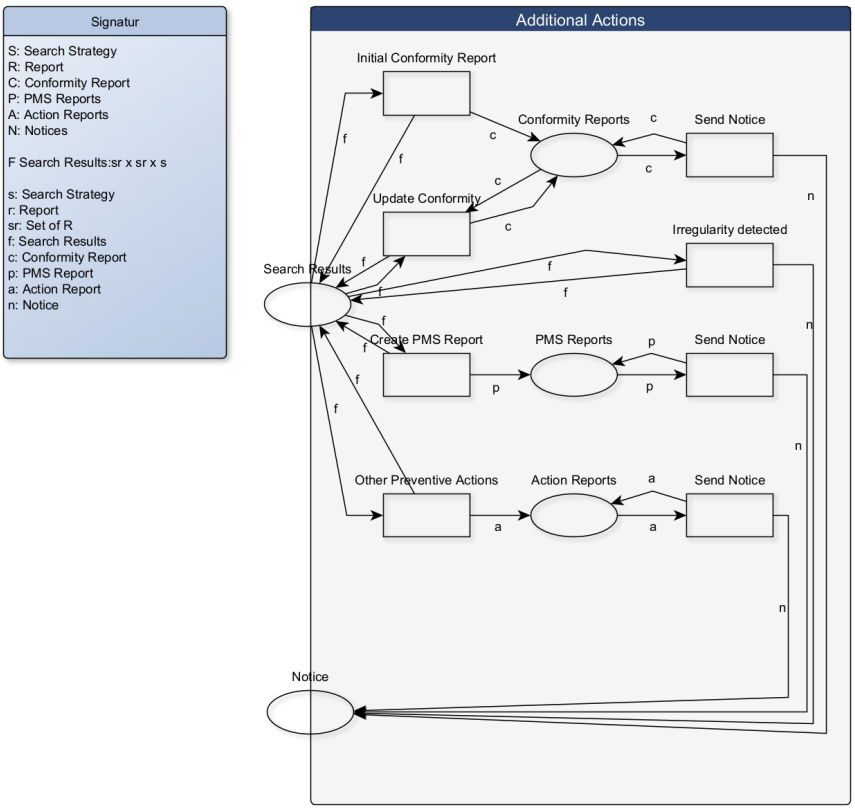

Figure 7: The inner of module *Additional Actions*

Strategy) and starts the data collection in module *Perform*. The submodule *Perform* consists of modules to retrieve and evaluate reports. Using the search profile, reports are retrieved in module *Query Data Sources* (Figure 8). Next, the query results are evaluated in module *Assess Query Results* (Figure 9). In the end, all relevant information are collected and passed as search results to module *Additional Actions* (Figure 7). In case the search strategy does not give the desired result, it can be revised in module *Modify Search Strategy* which also triggers to start the whole module *Perform* from scratch using a *Modified Search Strategy*.

An example how the inner of a module looks like is given in Figure 7 which shows module *Additional Actions*. The signature on the left defines the data and its structure. For instance, search results F is a tuple of the relevant search results, the not relevant search results and the search strategy used to query the databases. Using the search results gained in module *Data Collection* several actions are performed in this module. For instance, an initial conformity report (denoted as c) can be created based on the results, an existing one updated, a PMS report created or other preventive actions can be triggered. All actions produce a notice (denoted as n) that is returned to module *Devices in Market*.

3. Challenges and AI-Support

In the following, we present challenges regulatory experts face and propose suitable techniques how AI-support for such situations can be given on two modules, *Query Data Sources* and *Assess Query Results*. Red transitions indicate that support can be given.

3.1. Querying all relevant Data Sources

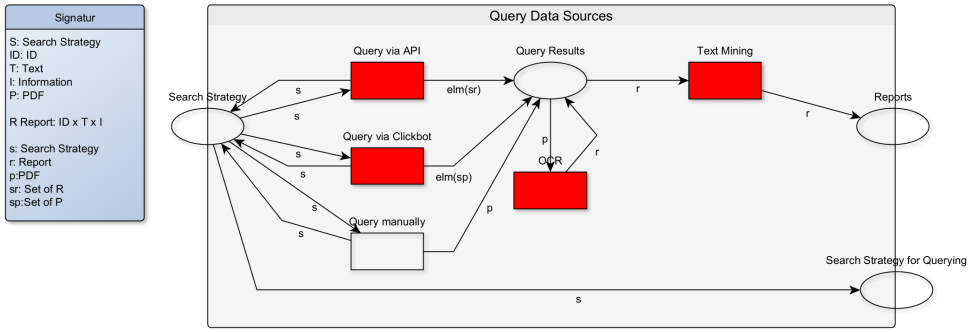

Figure 8: The inner of module *Query Data Sources*.

Once the search profile is defined, reports from different databases have to be retrieved. There are diverse data sources with different access options and data schemes. While some data sources can only be accessed manually and do not provide a search function, others offer the possibility to retrieve reports via API using keywords and filters. For example, the BfArM database is a mandatory data source for PMS in Germany containing warnings and notes from medical device manufacturers. However, the portal does neither offer the possibility to query the reports via API nor to search for specific reports. Instead, all released documents can only be viewed chronologically. Furthermore, documents need to be downloaded manually as PDFs in a non-standardized format. On the other side, Medline offers an extensive API that allows users to query the database directly using keywords and filters returning the reports in a machine-readable and structured format (XML). Figure 8 shows the module *Query Data sources* in detail. The module contains the mentioned steps to retrieve reports from relevant data sources and prepare them for further analysis. As retrieving the reports from the data sources is a repetitive and time consuming task, we present AI-techniques supporting this part.

3.1.1. AI-support for *Query Data Sources*

Depending on the type of data source, data can either be queried via API resulting in machine-readable reports (denoted as *r*), via click-bots that query the database as a human would do or manually, yielding PDFs (denoted as *p*). Automatically querying via API or click-bots gives a list or set of reports or PDFs (denoted as *sr* or *sp*) which are element-wise put into the place *Query Results* using the *elm* function [4]. The following steps, *OCR* and *Text Mining*, take a PDF or report as input and output one report again. For simplicity, a report *r* contains a report

identifier, a text, i.e., its content and additional information. Additional information can be additional fields of information gained when querying the report or structured information added by processing steps in *OCR* or *Text Mining*.

Automatically querying BfArM can be done using a simple click-bot which downloads all PDFs. Optical Character Recognition (OCR) and techniques from document understanding are applied to extract the text of PDFs in a structured format [6] and convert it into a report. Having the reports in an electronically readable format, text mining can be applied to harmonize and standardize the reports, e.g., through entity tagging and automatic extraction and mapping of text fields [7]. Using these techniques allows to query and extract reports and text from different databases in a productive manner with the result of having all reports in a structured and harmonized format for further analysis.

3.2. Assessing the Query Results

When searching in databases, often more reports are retrieved than are actually relevant. For instance, Medline contains mainly clinical studies which can contain adverse events. Unfortunately, Medline does not index adverse events and offers no possibility to filter by them. However, finding the relevant reports or information within reports is a crucial step. The next module deals with the assessment of the reports retrieved by the module *Query Data Sources*. For a report to be relevant, there mainly are two qualitative reasons as mentioned by the regulatory specialists:

- The report contains an unknown problem.
- The report contains a known problem but with an unknown cause.

Depending on the type of search being conducted, there might be other reasons why a report is relevant. Furthermore, quantitative and statistical information from the query can also be a piece of relevant information, e.g., if a rising number of known or unknown problems with a certain medical device over a period can be noted. This can be found by investigating the collection of search results.

Regulatory specialists tend to reduce the number of retrieved reports by experimenting with different search strings and using filter until an appropriate number of reports is reached. Afterwards, each report is examined to collect the relevant information within the report. Based on the relevant information within the report, each report is classified into different categories. Basically, a report is either relevant or not relevant. However, as MDR requires to justify why reports are classified as relevant or not relevant, some companies use finer-grained classification schemes. For instance, they classify scientific publications according to their level of evidence, i.e. their quality. The level of evidence is then used as one parameter for justifying if a report is relevant or not. For example, for clinical trials, which are a type of scientific publications, the evidence is determined based on the number of patients, interventions taken and the outcome. This information is referred to as PICO-Parameters [8] which is also a procedure to assess reports and collect information. Other relevance parameters might be the appropriate population group, data quality or overall study quality. In addition, when finding reports dealing with a similar devices, it must be investigated whether this also applies for the own device.

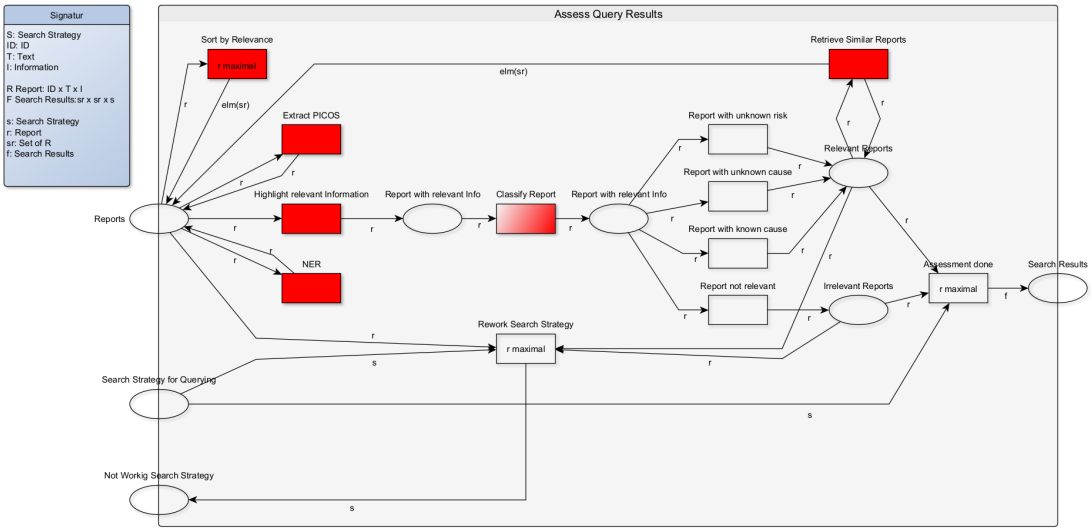

Figure 9: The inner of module *Assess Query Results*.

Module *Assess Query Results* in Figure 9 shows how reports are assessed. In case the user finds that the search strategy used produces not the desired result, e.g. if too many reports have been retrieved, the search strategy used is passed back to module *Modify Search Strategy*. Similarly, the strategy is passed back if the relevance of the reports after the assessment is too low. This starts the querying part again.

3.2.1. AI-support for *Assess Query Results*

Sorting the reports The retrieved reports are typically not ordered by their relevance with respect to a specific company. Given the reports retrieved from the query, the reports should be ordered so that the most relevant reports are shown first and the least relevant ones at the end. This happens in *Sort by Relevance*. It takes all reports r and adds a relevance score to be used for sorting the reports.

The assessment of relevance with respect to the query can be done automatically using methods like inverse term frequency and field length (TF-IDF). One popular TF-IDF function is BM25 [9]. More recently, neural ranking methods have been proposed that use neural networks to assess the relevance of a report. For instance, based on the number of clicks a report receives given a certain query [10]. Such methods can also be customized by assessing the report's content and taking previous assessments regarding the same device into account. Techniques to realize a customized ranker are, for instance, called Learning-2-Rank (L2R) [11] which are based on neural ranking models. Such models are also able to rank based on different relevance levels, e.g. for different types of impact a recently marked report had.

Extracting relevant information Finding and extracting the relevant information is an important step before being able to classify the reports. Certain entities of interest can be

extracted using techniques such as Named-Entity-Recognition [7]. Parameters such as PICO can be extracted from clinical studies automatically [8]. Furthermore, a similar technique can be used to automatically assess the evidence of a clinical study. For scientific publications, it is possible to determine if the publication is a randomized clinical trial [12]. Using the decisions from previous assessments, it is also possible to learn what parts of a report contain the information that is of interest [13]. In order to support the user to quickly decide about the relevance, extracted information can be highlighted to make the human pay attention to the parts that may be of interest. All these methods work automatically and add or highlight information as additional information to a report.

Report classification By using the extracted information in the reports, it can be assessed and classified according to the individual classification scheme of a company. While the AI can support the user by proposing a classification, the actual decision should always be made by the human. Thus, *Classify reports* is not colored completely red. Using information from recent assessments, we have seen that we are able to train a ML-algorithm to propose a classification to the user.

Retrieval Given a specific report which was ranked as relevant, the user might be interested in reports that are very similar to that one. For this, all reports are compared against the report of interest regarding a certain metric such as the cosine distance from features produced, e.g., by TF-IDF obtaining a set of new reports that can be assessed. This allows the user to quickly find reports that are of specific interest. Those reports can also be retrieved from data sources that were not queried originally.

Customized AI Given the continuous nature of PMS within the lifecycle of a medical device, historic data from previous assessments is likely to be available, i.e. data which reports and passages are relevant and which are irrelevant for a given medical product from a certain company. Using this information, the information extraction, relevance assessment and retrieval can be customized to the specific needs of a customer or company.

4. Discussion & Conclusion

In this work, we have proposed a way for structuring and modelling the activities supporting the safety and conformity of medical devices as a Heraklit scheme. The processes and the scheme are based and approved in interviews with regulatory specialists. Furthermore, AI-support along the process has been shown.

While the scheme supports how three companies perform those processes, we cannot claim completeness to support all possible ways. Furthermore, the processes have yet not been used for creating reports that are accepted by notified bodies. Furthermore, a more extensive evaluation would be required to further assess the scheme's usefulness, correctness and conformity, also with respect to other companies and other regulations.

However, the scheme gives an overview of relevant activities and details of important parts. Due to the modularization, it can be modified to specific needs or requirements.

In the future, we plan to use the scheme for developing information systems that support humans in PMS, Vigilance and related activities. Whenever possible, AI-solutions will be added that take over the repetitive parts of this work, i.e., help in collection and analysing reports and support in finding the relevant information within the large amount of information available. Special attention will be payed to regulatory requirements. Furthermore, we will evaluate the applicability and usefulness of the information system, which is based on the scheme, in supporting regulatory specialists in their daily work. This also involves an evaluation if the reports created using the tool will be accepted by notified bodies.

Acknowledgments

Part of this work has been done in the project *SmartVigilance* (FKZ: 01IS20028C) with financial support by the Federal Ministry of Education and Research (BMBF). We would like to thank pfm medical ag (https://www.pfmmedical.de/) and AkknaTek (https://www.akknatek.com/) for their support.

References

[1] Regulation (EU) 2017/745 of the European Parliament and of the council of 5 april 2017 on medical devices, amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and repealing Council Directives 90/385/EEC and 93/42/EEC, 2017.

[2] J. Harer, C. Baumgartner, Anforderungen an Medizinprodukte: Praxisleitfaden für Hersteller und Zulieferer, Carl Hanser Verlag GmbH Co KG, 2021.

[3] W. Reisig, Understanding petri nets: modeling techniques, analysis methods, case studies, Springer, 2013.

[4] P. Fettke, W. Reisig, Modelling service-oriented systems and cloud services with HERAKLIT, Advances in Service-Oriented and Cloud Computing, Springer International Publishing, 2021, pp. 77–89. doi:10.1007/978-3-030-71906-7_7.

[5] J.-R. Rehse, P. Fettke, A procedure model for situational reference model mining, Enterprise Modelling and Information Systems Architectures (EMISAJ) – International Journal of Conceptual Modeling 14 (2019).

[6] D. Baviskar, S. Ahirrao, V. Potdar, K. Kotecha, Efficient automated processing of the unstructured documents using artificial intelligence: A systematic literature review and future directions, IEEE Access 9 (2021) 72894–72936. doi:10.1109/ACCESS.2021.3072900.

[7] W. W. M. Fleuren, W. Alkema, Application of text mining in the biomedical domain, Methods 74 (2015) 97–106. doi:https://doi.org/10.1016/j.ymeth.2015.01.015.

[8] B. C. Wallace, J. Kuiper, A. Sharma, M. Zhu, I. J. Marshall, Extracting PICO sentences from clinical trial reports using supervised distant supervision, J. Mach. Learn. Res. 17 (2016) 4572–4596.

[9] S. Robertson, H. Zaragoza, The probabilistic relevance framework: BM25 and beyond, Now Publishers Inc, 2009.

[10] S. Mohan, N. Fiorini, S. Kim, Z. Lu, A fast deep learning model for textual relevance in

biomedical information retrieval, in: Proceedings of the 2018 World Wide Web Conference, 2018, pp. 77–86.

[11] B. Mitra, N. Craswell, Neural models for information retrieval, arXiv preprint arXiv:1705.01509 (2017).

[12] I. J. Marshall, A. Noel-Storr, J. Kuiper, J. Thomas, B. C. Wallace, Machine learning for identifying randomized controlled trials: An evaluation and practitioner's guide, Research Synthesis Methods 9 (2018) 602–614. doi:https://doi.org/10.1002/jrsm.1287.

[13] W. Yang, H. Zhang, J. Lin, Simple applications of BERT for ad hoc document retrieval, arXiv preprint arXiv:1903.10972 (2019).