








Verifying Guideline Compliance in Clinical Treatment Using Multi-perspective Conformance Checking: A Case Study

Joscha Grüger^{1,2} , Tobias Geyer² , Martin Kuhn² ,
Stephan A. Braun^{3,4} , and Ralph Bergmann^{1,2} 

¹ Business Information Systems II, University of Trier, Trier, Germany
grueger@uni-trier.de

² German Research Center for Artificial Intelligence (DFKI), Trier, Germany

³ Department of Dermatology, University Hospital Münster, Münster, Germany

⁴ Department of Dermatology, Medical Faculty, Heinrich Heine University,
Düsseldorf, Germany

Abstract. Clinical guidelines support physicians in the evidence-based treatment of patients. The technical verification of guideline compliance is not trivial, since guideline knowledge is usually represented textually and none of the approaches to computer-interpretable guideline representation has yet been able to establish itself. Due to the procedural nature of treatment sequences, this case study examines the applicability of a guideline process model to real hospital data for verification of guideline compliance. For this purpose, the limitations and challenges in the transformation of clinical data into an event log and in the application of conformance checking to align the data with the guideline reference model are investigated. As a data set, we use treatment data of skin tumor patients from a cancer registry enriched by hospital information system data. The results show the difficulty of applying process mining to medically complex and heterogeneous data and the need for complex preprocessing. The variability of clinical processes makes the application of global conformance checking algorithms challenging. In addition, the work shows the semantic weakness of the alignments and the need for new semantically sensitive approaches.

Keywords: Process mining · Multi-perspective Conformance checking · Clinical guidelines · Guideline compliance

1 Introduction

Evidence-based medicine states that patient-centered medical treatment decisions should be based on empirically proven effectiveness whenever possible [23].

The healthcAIre project is funded by the ministry of science and health of the German state Rhineland-Palatinate and the Pre-OnkoCase project is funded by the National Care Conference on Skin Cancer (NVKH) e.V.

© The Author(s) 2022

J. Munoz-Gama and X. Lu (Eds.): ICPM 2021 Workshops, LNBP 433, pp. 301–313, 2022.

https://doi.org/10.1007/978-3-030-98581-3_22

This knowledge is documented in clinical guidelines [25]. The degree to which clinical treatment processes in practice are guideline-compliant and thereby evidence-based is unknown [9, 14]. The verification of guideline compliance is relevant, e.g., for the certification of oncology centers¹, the development of clinical decision support systems [2, 15] and medical research [3, 12]. An approach to check the compliance of treatment processes against guidelines is to interpret individual treatment processes as process instances and the guideline as a reference process model [10]. This would enable the use of conformance checking, a process mining technique, which raises two challenges though. First, the transformation of the guideline knowledge into a process model. Second, the provision of clinical data as event logs and their preprocessing. Due to the lack of standardization of clinical data storage and the associated structure heterogeneity, naming and data quality, a preprocessing of this data is necessary [13]. Furthermore, clinical processes are characterized as highly variable, ad-hoc, multidisciplinary and vary from hospital to hospital [22].

As part of the Pre-OnkoCase² project, a process model was developed for a section of the malignant melanoma guideline. In this case study, we investigate to what extent the model can be applied to real clinical data, what preprocessing is necessary and what limitations exist. The case study was conducted in collaboration with the skin tumor center of the Münster University Hospital³.

The remainder of the paper is organized as follows. Section 2 provides background information about the medical context, the process reference model and conformance checking details. Section 3 describes the research method and shows how the event log is created. Section 4 describes the implementation of the conformance checking and the required preprocessing. In Sect. 5, the results are discussed and Sect. 6 concludes the paper.

2 Background

Within the Pre-OnkoCase research project, we investigated how a clinical guideline can be represented procedurally. Since guidelines assume tacit knowledge, they provide an incomplete representation of the treatment processes. Therefore, missing information had to be supplemented by experts' knowledge. In workshops with domain experts of the skin tumor center in Münster, a conceptual model of a section of the evidence-based guideline for the treatment of malignant melanoma (skin cancer) [6] was created. Due to the size and complexity of the model Fig. 1 shows just a sketch of the fundamental treatment process. Each element of the sketch is a representative of a treatment section in the treatment courses of patients who have been diagnosed with melanoma and consists of a set of activities. If malignant melanoma is diagnosed during the clinical and histopathological examination, which is part of the *Diagnosis of Melanoma* section, then several treatment options are available to the patient.

¹ <https://www.krebsgesellschaft.de/deutsche-krebsgesellschaft/zertifizierung.html>.

² <https://nvkh.de/projekte/pre-oncocase>.

³ <https://www.ukm.de/index.php?id=hauttumorzentrum>.

- *Re-Excision*: Repeated excision ensures that no tumor residues remain.
- *Sentinel Lymph Node Biopsy*: The patient can receive a lymph node sonography and receives a re-excision together with the sentinel lymph node biopsy.
- *Other Diagnostic Measures*: The patient can receive diagnostics to confirm metastases or further examinations with imaging techniques.
- *Staging up to IIB & Staging IIC and III*: Depending on prior examinations, the patient can come to staging, which provides patients and physicians the critical benchmark for defining prognosis and for determining the best treatment approach [8].
- *Lymphadenectomy*: The patient receives a lymphadenectomy, can receive radiotherapy afterwards and receives a drug therapy.
- *Adjuvant Therapy*: The patient receives additional cancer treatment after initial treatment to reduce the risk of recurrence.




Fig. 1. Reduced overview of the considered clinical guideline section for the treatment of malignant melanoma patients. Paths marked with “...” are treatment areas, which are not covered in the model, e.g., the treatment of stage IV patients or follow-up care.

The final conceptual model was then transferred into a Data Petri Net (DPN) by Geyer [11]. A DPN is an extended Petri net that can map data and time information [16]. The modeled DPN consists of 50 places and 76 transitions (see Fig. 2). Due to the many decisions made in treatment based on examination results, there are 52 transitions with a guard. The resulting model represents all conditions and recommendations of the selected guideline section.

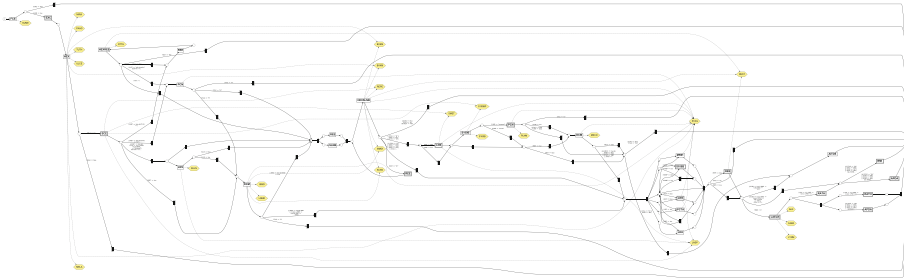


Fig. 2. Overview of the DPN reference model

In the following, the basic terminology in the context of multi-perspective conformance checking is explained. Multi-perspective conformance checking describes the process of identifying discrepancies between the desired behavior of the process, represented by the process model, and the actual behavior involving multiple perspectives such as the data perspective or the time perspective. Most approaches use alignments for this purpose, which are a mapping of the process instance to the process model. In the context of alignments, a *log move* is executed by the alignment algorithm for events that are recorded in the event log but do not occur in the process model. A *model move* is executed if events occur in the process model but do not occur in the event log. If the event from the event log matches the activity in the process model, but the values of the variables do not match, this is called *incorrect synchronous move*. If everything matches, the move is defined as *correct synchronous move* [16].

Most of the process mining algorithms which are capable of calculating multi-perspective alignments are using the Alpha* algorithm [5] in combination with MILP (Mixed Integer Linear Programming) [24]. The state-of-the-art approach is from Mannhardt [17] where DPNs are used for calculating multi-perspective optimal alignments [7]. It is also possible to calculate multi-perspective alignments by using MP-Declare a multi-perspective version of Declare [21]. This method was developed by Mawoko [19] and utilized a similar approach as Mannhardt.

3 Research Method

The exemplary data set used in this project represents the treatment of a total of five real patients diagnosed with malignant melanoma from Münster University Hospital. For data privacy reasons, the data were anonymized. The treatment data are provided in the format of the ADT/GEKID basic data set⁴. The uniform oncological ADT/GEKID basic data set describes a common coding scheme for the documentation of oncological treatments in Germany in the form of an XML schema. A major advantage of using data in the format of the basic data set is that it is used by all German cancer registries and results are thus transferable

⁴ <https://www.gekid.de/adt-gekid-basisdatensatz>.

and comparable. The basic data set includes among others patient master data, diagnostic data, histology data, cancer classification data, surgical data, therapy data and tumor conference data.

Each entry in the basic data set is provided with a timestamp and a treating resource and uniquely assignable to a patient and a treatment case. The structure of the basic data set is based on the obligation of hospitals in Germany to report the course of cancer cases to cancer registries. Accordingly, the data on individual treatment activities are assigned to reporting elements in the XML format and enriched with treatment-specific information. In order to apply conformance checking, the data are transferred into the XES event log format [1]. For this purpose, a generic XML to XES converter was implemented in Python and configured to convert ADT/GEKID data to XES.

The resulting process log covers many areas important for determining guideline compliance, such as surgeries and diagnoses. Also, additional information on follow-up examinations, medical therapies and tumor conferences are contained. However, it lacks information on, e.g., histological examinations, certain tumor markers, or lymphadenectomy. The resulting event log contains 24 different events while considering different medical procedures as different events.




Fig. 3. Overview of the process steps up to conformance checking

In order to be able to take these data into account in conformance checking, the log was enriched with treatment data from the hospital information system (HIS). For this purpose, data from the HIS were exported as CSV and imported into the XES file. Most of the entries could be transferred automatically, since they are structured and timestamped. However, individual details of the treatment process had to be extracted manually from the free text of the diagnostic findings and doctor’s letters. The final event log contains 179 different events and a total of 1114 events, an average of 222 events per patient.

4 Implementation

The following describes the adjustments that were necessary to perform conformance checking. An overview of the individual procedures in the project are shown in Fig. 3.

4.1 Preprocessing

The final event log contains 179 different events, while the guideline reference model has only 20 different events. The difference results from the fact that the event log contains events of other medical domains such as nursing and psychosocial care and from the fact that the granularity in which events are represented is inconsistent. In addition, it is also due to the fact that there are deviations from the guideline. In order to perform an alignment between the event log and the guideline reference model, extensive preprocessing had to be performed: removal of explicitly irrelevant events, reduction of therapy events to the respective initial therapy event, harmonization of granularity, event aggregation and event and variable name matching.

In the first step, events that were explicitly irrelevant for conformance checking were removed. These include events from perspectives not considered by the guideline, such as the nursing and psychosocial domains, events such as tumor conferences, which neither establish new diagnoses nor provide direct treatment, and events such as follow-up care, which are outside the selected guideline section. The events were identified using the event names and a HIS-internal ID. Subsequently, in the second step, the therapy sequences of the same therapy were reduced to the respective starting event. This is necessary because cancer therapies are usually performed several times and the reference model of the guideline, however, only addresses whether a patient with certain diagnoses receives a certain therapy and then implies that this therapy is subsequently performed correctly. The granularity of the event log in terms of the event description is in many ways finer than in the reference model. While the reference model refers to “excision”, the ICPM (International Classification of Procedures in Medicine) classification used in the data set defines over 30 different excisions. Therefore, the data set is harmonized in terms of granularity. For this purpose, the ICPM code is abstracted in the hierarchically structured coding scheme to such an extent that the description matches the identifiers of the reference model. This results in partial events with identical designation and identical timestamp, which originally described, e.g., surgeries with several similar individual events are aggregated to one event. It is essential that the names of the same variables and events in the guideline reference model and in the event log are identical. For this purpose, a comparison of the identifiers of the event log with identifiers of the model was performed. This was particularly time-consuming because identifiers were not consistent and unique. This is on the one hand due to the fact that the data set is based on data from two systems and on the other hand due to the fact that the treatment documentation is partially in free text and identifiers were accordingly heterogeneous. The resulting event log forms the basis for the conformance check. After applying the described preprocessing steps, the event log only contains 40 different events directly related to treatment instead of the initial 179. The discrepancy between 20 activities in the model and 40 events in the log was deliberately accepted in order to have complete traces and a comparison between guideline specifications and reality.

4.2 Conformance Checking

The in the following presented conformance checking approach is considered as a *global conformance checking* technique, which views the process reference model as an accurate representation of the overall process behavior. It is assumed that the whole process is modeled and can therefore be checked. This method enables not only the identification of the deviations but also the identification of the exact source causing the problem [17]. We have chosen a global conformance checking approach as this corresponds to the medical practice of considering entire treatment processes.

Therefore, ProM [26] was used with the Multi-perspective Process Explorer (MPE) [18], which uses the multi-perspective alignment algorithm developed by Mannhardt [17]. A fundamental feature of the conformance checking algorithm is the definition of a cost function. The cost function should be defined in a way that the calculated alignments are semantically correct. We define semantically correct alignments as a meaningful and logical alignment concerning a process instance with deviations. A semantically correct alignment does not need to be an optimal alignment but should be correct in a sense that a domain expert would consider this alignment as meaningful.

First, the standard cost function is used. This cost function defines the cost as 3 for log move (delete), 2 for model move (insert), and 1 for incorrect synchronous move (data write). The standard cost function results in semantically incorrect alignments, because the alignment algorithm changes the attribute values of events to create an optimal alignment. In the medical context, this is semantically incorrect, as the data collected by the doctor represents reality and should be immutable for the algorithm. In this case, the standard cost function generates unusable alignments.

To achieve the desired result, the cost for data writes is increased such that it is higher than for the other two operations. Also, the delete cost for events that are not part of the staging process is reduced to 0, since in the course of the medical examination it is possible that multiple additional examinations are undertaken, that are needed to perform but are not depicted by the process model. Thus, costs were defined as 1 for log move (delete), and model move (insert), 0 for log move (for events not defined in the model) and incorrect synchronous move (non-data write) and 2 for incorrect synchronous move (data write).

Moreover, it is important to mention that the cost for the non-data writes was set to 0, because this allows the alignment algorithm to make insert operations that are associated with attribute values, without paying the cost for the data writes. The calculated fitness value itself was not considered, since the focus of the use case is on the calculated deviations on the event level. Due to privacy regulations it is not possible to show the resulting alignments, thus the results will be explained in a qualitative way. Two of the alignments are semantically correct. These traces correspond perfectly to the guideline but also contain medical examinations, which are not depicted in the process model, but were needed to perform. The additional undertaken examinations are deleted by the

alignment algorithm, which is semantically correct since these examinations are a positive deviation from the guideline, which have the cost of 0. For the other three traces a semantically correct alignment was not possible. This is mainly due to the occurrence of events that are depicted of the process model but are occurring at other positions as expected. In this case, the alignment algorithm seeks the shortest or least expensive path through the process model and deletes correct events or inserts new events, which already have been executed. Here, the least expensive path allowed by the process model is not semantically correct in every case, as our alignments show. As stated by [17] the algorithm only shows one alignment and this is the optimal one in terms of alignment costs. However, there are also possible alignments that might be better in terms of semantics but worse in terms of alignment costs.

In summary, it was not possible to define a cost function which leads to semantically correct alignments for all traces. Nevertheless, it was possible to identify the medical examinations that are not part of the guideline but were executed by the physician.

5 Discussion

The following section discusses the results of the project and the associated problems and limitations identified. Although the domain experts attempted to provide the most heterogeneous and complex patient data possible for this case study, it should be noted that additional challenges and issues may arise as additional patient cases are examined.

Several problems, partially typical for medical data, were found in the data set used. The following issues and characteristics were identified: high variability of treatment processes, time delays, incomplete data, none-activity-data and mapping ambiguities between reference activities. The treatment histories have a **high degree of variability** typical for medical data. Patient treatment data have shown that there are activities in treatment that can occur at any time and any number of times. Thus, such activities occur more frequently than described in the guideline. These treatment activities pose a challenge in guideline compliance checking because they are explicitly mentioned in the guideline only at specific points in treatment. Consequently, guideline-compliant modeling does not represent all contingencies of medical treatment, leading to the identification of activities as deviations where they do occur additionally. Another important aspect at this point is that some of these activities may play a crucial role in the further course of treatment. For this reason, the activities must be able to occur at any time in the model and they must have paths to all possible subsequent treatments. However, based on the data collected so far, it is apparent that mapping all options would increase the complexity of the model and thus the effort to maintain it is no longer manageable.

A similar problem occurs due to **time delays** in treatment. For example, in the treatment of patients, surgical procedures are followed by histological laboratory examinations in which, e.g., tissue or lymph nodes are examined. Consequently, the obvious modeling approach is to place the laboratory testing after

surgery. In practice, treatment data have shown that some time elapses between surgery and laboratory examination, and patients continue to receive treatment in the meantime. This leads to the issue of valid activities being identified as a deviation or violation. The data from the systems are **incomplete** as they only represent the clinically documented course and parts of the out-of-hospital treatment and diagnosis are missing. This is particularly evident in the data for events at the beginning of the treatment process. Although it is evident from a medical view that all patients should have passed through the same diagnostic steps, patients start with different events. This is due to the fact that parts of the treatment such as excision, histological examination, initial clinical examination, etc. were performed out-of-hospital. Parts of the ADT/GEKID data are **none-activity-data** and thus cannot be assigned to an event or timestamp. For this data, it is neither possible to determine when nor in the course of which activity it was collected. This affects the master data, which also includes attributes such as age, which are crucial for guideline recommendations. The same applies to the diagnostic data, which only reflects the current status and not the procedural progression over time. Therefore, it is not possible to track staging over the progression of treatment with ADT/GEKID. In the context of the reference model, **mapping ambiguities between reference activities** occurs in the data. Thus, there are events in the event log which could imply the execution of certain activities by numerous attributes. However, the collection of the value does not necessarily imply the use of the value and thus the execution of the activity in the process model. Standard laboratory tests, e.g., involve the collection of numerous values, including tumor markers. However, the documentation of the values does not allow any conclusion to be drawn about the observation, analysis and usage of the tumor markers. Thus, at no point in the process can it be determined whether a particular tumor marker was considered or not.

During **conformance checking**, process mining specific problems were identified in addition to the data set related ones. The following problems have been identified: semantically inappropriate control-flow alignments, semantically inappropriate data alignments and definition of cost function. The **semantically inappropriate control-flow alignments** describe a conflict between the goal of the algorithm and the medical intent. By default, the algorithm uses a cost function where aligning data values is cheaper than aligning events. As a result, patient examination values are modified during alignments, such as changing the staging value, to restore conformance. The examination values are of utmost importance for the course of treatment, but should only be modifiable by new diagnoses of the physicians and not by the algorithm. Accordingly, to produce the desired behavior, in the configuration, aligning data values is more expensive than aligning events. As a result, the sequence of events is aligned, but not in the desired way. Consequently, situations arise where the alignment changes only a single data value, e.g., making it the most favorable path for the alignment. This approach ends the patient's path as fast as possible and implies, e.g., that no melanoma was found during the initial clinical examination and the patient is discharged from the hospital. Therefore, from a medical point of view,

it becomes apparent that the most favorable path represents not the best possible course of treatment. Based on this finding, further efforts should be made to examine whether the current conformance checking approach is suitable for checking medical treatment processes for guideline compliance. Since treatment courses are highly dynamic, a potentially more appropriate approach would be to examine whether a possible path of the process model can be reconstructed via sequence segments of the corresponding treatment course. Since guideline specifications only partially describe steps or sequences anyway, an alignment of sequence segments would provide a means for medical conformance checking. A suitable approach could be a *local conformance checking* technique, which describes the process of checking the conformance by using a set of independent rules regarding the process. Therefore, only specific parts of the process are checked not the process as a whole. These rules are often defined in LTL or in declarative modeling languages like Declare [4]. Furthermore, **semantically inappropriate data alignments** could be identified when performing conformance checking. These occurred when a guard was violated by an improper value. For example, a patient may receive radiotherapy after a lymphadenectomy if they have a count of three or more lymph nodes affected with cancer. In an alignment, the value was generically set to 1000, which satisfies the condition but creates semantic incorrectness. At this point, it becomes evident once again that data values should not be adaptable across the board in medical conformance checking. The medical context is highly relevant in and between treatment steps, which is why simple value alignments to satisfy guards are not sufficient. If a conformance checking algorithm should indeed have the authority to make data alignments, then semantic technologies must be used in order to draw proper conclusions and achieve meaningful results. Another problem became apparent in the attempt to define a generally valid **cost function** for the patients. Thus, although desired alignments could be achieved sporadically by changing costs, they could only ever be achieved for an individual patient. Since the medical conditions for a patient in treatment are highly dynamic and individual, it is not possible to achieve globally desired results by defining costs.

6 Conclusion

In this work, we focused on the applicability of conformance checking to determine clinical guideline compliance on clinical data. For our case study, we used real data of non-trivial treatment and diagnosis of malignant melanoma provided by Münster University Hospital and a procedural guideline representation created in collaboration with medical professionals. The data used were in the format of the ADT/GEKID data set, which is used by the German cancer registries, and enriched with data from a HIS where necessary.

We showed that it is possible to use conformance checking to verify clinical guideline conformance of real-world clinical data. Unfortunately, there are a number of application problems, mostly rooted in the data, but also in the conformance checking algorithm and the process model. In particular, the characteristically high variability of clinical treatment processes is a challenge. Both

the execution and the order of execution of activities in clinical treatment processes are subject to a variety of factors, including co-morbidities, time delays in the process and patient preferences, resulting in highly variable processes. In addition, incomplete processes, e.g., when data from treatments in other organizations are not available, need to be handled. Moreover, alignments by writing attribute values or deleting activities partially resulted in semantically incorrect alignments. Further challenges lie in the preprocessing of the data, as they were inconsistent in granularity, contained activities irrelevant to conformance checking, and most importantly were documented heterogeneously and partially unstructured, requiring a complex preprocessing process.

We plan to extend the evaluation to other guidelines, including time-constraints such as follow-up care. We are also working on fitness functions based on sub-processes and an analogy-based alignment approach. In this context, we plan to further investigate the clinical data and define similarity measures for treatment-relevant parameters with medical experts. Also, we want to test other approaches such as deep-align [20] and investigate how they address the identified problems. Many problems are due to semantic violations of the alignment. Here, we are working on an ontology-supported hybrid alignment procedure that detects semantically incorrect alignments and tries to prevent them.

References

1. IEEE Standard for eXtensible Event Stream (XES) for Achieving Interoperability in Event Logs and Event Streams. IEEE Std 1849–2016, pp. 1–50 (2016)
2. Alharbi, R.F., Berri, J., El-Masri, S.: Ontology based clinical decision support system for diabetes diagnostic. In: Science and Information Conference (SAI), pp. 597–602. IEEE (2015)
3. American Society of Clinical Oncology: Good clinical practice research guidelines reviewed, emphasis given to responsibilities of investigators: second article in a series. *J. Oncol. Pract.* **4**(5), 233–235 (2008)
4. Caron, F.: Business process analytics for enterprise risk management and auditing (2013)
5. Dechter, R., Pearl, J.: Generalized best-first search strategies and the optimality of a*. *J. ACM* **32**(3), 505–536 (1985)
6. Deutsche Krebsgesellschaft, Deutsche Krebshilfe, AWMF: Leitlinienprogramm Onkologie: Diagnostik, Therapie und Nachsorge des Melanoms, Kurzversion 3.3, 2020f (30042018). <https://www.leitlinienprogramm-onkologie.de/leitlinien/melanom/>. Accessed 28 May 2021
7. Dunzer, S., Stierle, M., Matzner, M., Baier, S.: Conformance checking: a state-of-the-art literature review. In: Proceedings of the 11th International Conference on Subject-Oriented Business Process Management, pp. 1–10. ACM (2019)
8. Edge, S.B.: *AJCC Cancer Staging Manual*, 7th edn. Springer, New York (2010). <https://doi.org/10.1007/978-1-4757-3656-4>
9. Forsner, T., Hansson, J., Brommels, M., Wistedt, A.A., Forsell, Y.: Implementing clinical guidelines in psychiatry: a qualitative study of perceived facilitators and barriers. *BMC Psychiatry* **10**, 8 (2010)

10. Gatta, R., et al.: Clinical guidelines: a crossroad of many research areas. Challenges and opportunities in process mining for healthcare. In: Di Francescomarino, C., Dijkman, R., Zdun, U. (eds.) BPM 2019. LNBP, vol. 362, pp. 545–556. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-37453-2_44
11. Geyer, T.: Process mining on CIG process model representations for clinical guidelines. Master thesis, University of Trier (2021)
12. Grimshaw, J.M., Russell, I.T.: Effect of clinical guidelines on medical practice: a systematic review of rigorous evaluations. *Lancet* **342**(8883), 1317–1322 (1993)
13. Holzinger, A., Dehmer, M., Jurisica, I.: Knowledge discovery and interactive data mining in bioinformatics-state-of-the-art, future challenges and research directions. *BMC Bioinform.* **15**(Suppl 6), II (2014)
14. Landfeldt, E.: Compliance to care guidelines for Duchenne muscular dystrophy. *J. Neuromuscul. Dis.* **2**(1), 63–72 (2015)
15. Lenz, R., Reichert, M.: It support for healthcare processes - premises, challenges, perspectives. *Data Knowl. Eng.* **61**(1), 39–58 (2007)
16. de Leoni, M., van der Aalst, W.M.P.: Aligning event logs and process models for multi-perspective conformance checking: an approach based on integer linear programming. In: Daniel, F., Wang, J., Weber, B. (eds.) BPM 2013. LNCS, vol. 8094, pp. 113–129. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40176-3_10
17. Mannhardt, F.: Multi-perspective process mining. Ph.D. thesis, Technische Universiteit Eindhoven (2018)
18. Mannhardt, F., De Leoni, M., Reijers, H.A.: The multi-perspective process explorer. *BPM (Demos)* **1418**, 130–134 (2015)
19. Mawoko, C.T.: Aligning data-aware declarative process models and event logs. Master's thesis, University Tartu (2019)
20. Nolle, T., Seeliger, A., Thoma, N., Mühlhäuser, M.: DeepAlign: alignment-based process anomaly correction using recurrent neural networks. In: Dustdar, S., Yu, E., Salinesi, C., Rieu, D., Pant, V. (eds.) CAiSE 2020. LNCS, vol. 12127, pp. 319–333. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-49435-3_20
21. Pesic, M., Schonenberg, H., van der Aalst, W.M.: DECLARE: full support for loosely-structured processes. In: 11th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2007), p. 287 (2007)
22. Rebuge, Á., Ferreira, D.R.: Business process analysis in healthcare environments: a methodology based on process mining. *Inf. Syst.* **37**(2), 99–116 (2012)
23. Sackett, D.L.: Evidence-based medicine. *Semin. Perinatol.* **21**(1), 3–5 (1997)
24. Schrijver, A.: *Theory of Linear and Integer Programming*. Wiley, Chichester (1998)
25. Tobe, S.W., Hua, D., Twohig, P.: Clinical practice guidelines. In: Schiffrin, E.L., Touyz, R.M. (eds.) Hypertension, pp. 238–251. The Future Science Group eBook Collection, Future Medicine Ltd, London (2013)
26. van Dongen, B.F., de Medeiros, A.K.A., Verbeek, H.M.W., Weijters, A.J.M.M., van der Aalst, W.M.P.: The ProM framework: a new era in process mining tool support. In: Ciardo, G., Darondeau, P. (eds.) ICATPN 2005. LNCS, vol. 3536, pp. 444–454. Springer, Heidelberg (2005). https://doi.org/10.1007/11494744_25

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

