"The final publication is available at springerlink.com"
published at: Information Quality in e-Health
Lecture Notes in Computer Science Volume 7058, 2011, pp 533-546

Assessing Medical Treatment Compliance Based on Formal Process Modeling*

Reinhold Dunkl, Karl Anton Fröschl,
Wilfried Grossmann, and Stefanie Rinderle-Ma

University of Vienna, Austria
Faculty of Computer Science
{firstname.[midname.]lastname}@univie.ac.at

Abstract. The formalization and analysis of medical guidelines play an essential role in clinical practice nowadays. Due to their inexorably generic nature such guidelines leave room for different interpretation and implementation. Hence, it is desirable to understand this variability and its implications for patient treatment in practice. In this paper we propose an approach for comparing guideline-based treatment processes with empirical treatment processes. The methodology combines ideas from workflow modeling, process simulation, process mining, and statistical methods of evidence-based medicine. The applicability of the approach is illustrated based on the Cutaneous Melanoma use case.

Keywords: Healthcare Processes, Process Modeling, Process Mining

1 Introduction

Clinical practice is based upon medical knowledge, relating healing interventions causally to diseases, as well as upon clinical treatment routines integrating medical/pharmaceutical theory with best practices of patient care [5]. Medical guidelines, or standard operation procedures, codify the implementation of such clinical routines, in order to raise the specificity of patient care, to reduce the burden of medical decision making by defining applicable diagnostic criteria and intervention patterns, and to accumulate past experience in patient care as the state-of-the-art. In spite of their *prescriptive* intent, however, medical guidelines by necessity leave much room to adapt to specific circumstances and interpretations. Thus, while clinical evidence mirrors the conduct of patient care as defined through guidelines, empirical records of treatment rather reflect, in a *descriptive* way, the practical consequences of guideline implementation.

In what follows, the question is raised if there is anything to be learned from systematically contrasting medical evidence with guidelines and, as a prerequisite, which methodological frame, or toolbox, is required to support such a kind of comparison. To this end, it is hypothesized that, because of their inexorably

* The work presented in this paper was conducted in the context of the EBMC^2 project that is co-funded by the University of Vienna and the Medical University of Vienna.

generic nature, guidelines give rise to different ways and conditions of (compliant) implementation, the variability of which is becoming apparent only through the statistical analysis of formal reconstructions of clinical treatment processes. For one thing, it is to be expected that guideline implementation, on the level of the individual patient receiving a guideline-controlled treatment, undergoes manifold refinements accounting for the immanent variability of patients' health conditions, anamnestic peculiarities, adverse drug effects, and so forth.

In order to account for this empirical variability of patient treatment processes, a representation language providing sufficient level of descriptive detail, including diagnostics, interventions, medications, etc., in conjunction with a means to express the dynamics of treatment processes is essential. This requirement is met best by some formal process modeling language equipped with a temporal calculus [22, 3]. Depending on the respective domain of investigation, once the specific vocabulary of the medical domain is identified, individual clinical treatment processes can be restated formally as sentences of the process language, thus constituting the sample of process instances amenable to statistical scrutiny. On top of this symbolic reconstruction of empirical treatment processes, statistical analysis – more specifically, process mining methodology – is used to separate structural from (in principle) random components of process descriptions. At this stage, critical *structural* variations (if any) in guideline implementation are detected. Provided that medical guidelines themselves are restated in terms of formal process models (that is, as treatment process schemata), the formal difference between these "reference processes" and evidence-based process schemata becomes tractable. Furthermore, a suitable logic of comparison established between symbolically represented process schemata admits formal interpretation of differences in terms of a measure of (process) compliance. As yet, however, within the Evidence Based Medicin Compliance Cluster (EBMC^2) project, an *analytical model* is developed first, adapting and extending already existing simulation-based process mining proposals (for instance, such as [1]) with apt process distance measures and an exploratory mechanism for candidate substantiations of empirical process alterations; this conceptual approach is argued and presented subsequently.

Section 2 gives an overview of the proposed methodological framework and its reasoning, and links the proposal to related work in the field. In particular, the role and structure of formal modeling of medical treatment processes is introduced. Next, Section 3 turns to a practical illustration example, the case of Cutaneous Melanoma treatment, presenting a pertinent guideline in order to sketch approaches towards formal guideline representations. Based on Sections 2 and 3, Section 4 presents the main functional building blocks of the proposed framework used to simulate synthetic treatment processes on top of (i) formalized guideline-based process models, and (ii) patient samples. Furthermore, suitable process log data structures are considered, and data-analytical techniques of process aggregation and comparison, respectively, based on log data are discussed. The concluding section of the paper briefly discusses expected benefits of the framework, indicates how the conceptual approach is linked to pertinent data,

and points out still open issues in need of further research within the EBMC^2 project.

2 Methodologies of Modeling Clinical Routine

2.1 From Medical Evidence to Medical Guidelines

Medical evidence draws on empirical results about different treatments. PICO (cf. e.g., [21]: 113) is a standard scheme for obtaining such evidence from analyzing alternative therapies, prognosis of recovery/survival, diagnosis of health status, prevalence studies, etc. Using appropriate study designs, the PICO scheme produces, ideally, significant statistical evidence in terms of likelihoods, p-values, area under the ROC, sensitivity and specificity, or various risk measures. Methodological rigor of study design decides about the strength of evidence attainable (cf. [20, 26]), but yields rather isolated evidence:

- study focus is a single, well-defined question in order to obtain an accurate as possible answer;
- results are of a *statistical* nature, providing, at best, a numerical quantification of degrees of evidence over well-defined populations;
- integral treatment processes are dealt with in peacemeal fashion under standard conditions and, in particular, treatment *interactions* are barely addressed.

As a response, and remedy, medical *guidelines* seek to integrate the bits and pieces of medical evidence studies into coherent treatment processes in order to provide assistance to health professionals towards effective, high quality medical practice according to the best medical knowledge available from both, clinical research and expert consensus [10]. Formally, guidelines are composed of *key actions*, typically recommended conditionally depending on patients' sociodemographic and medical attributes (e.g., prior diagnoses), and usually indicating a degree of recommendation depending on the particularities of treatment instances and the medical setting. In general, key actions also carry some outcome measure impacting on future actions, perhaps supplemented with a quality measure.

Medical guidelines may be represented in varying degrees of formalization, from plain narrative to highly structured [25]. Workflows provide a rather natural means of expression, highlighting the "algorithmic" flavor of treatment processes and permitting a stepwise refinement of abstraction levels [23] including the annotation of branching and looping conditions. Fig. 1 sketches a top-level treatment flow pattern with building blocks to receive case-dependent iterative refinement (cf., e.g., Fig. 3 and 4 below).

Formal process modeling, however, has to account for a variety of flow conditions (such as unexpected treatment termination), the incidental interleave of concurrent treatment processes (e.g., in intensive care), as well as other contingencies, entailing discrepancies between ideal guideline implementation and

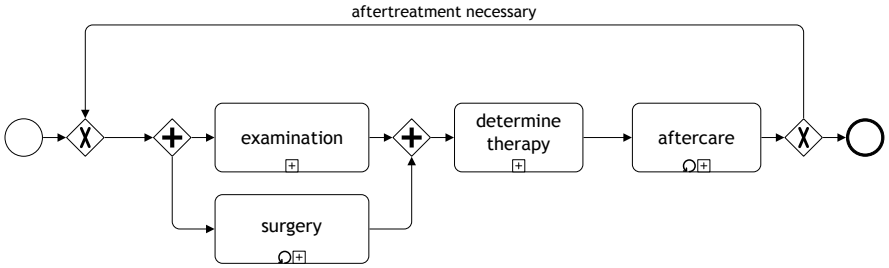

Fig. 1. A High-level Treatment Process Flow Representation

actual clinical care practice. This is reflected in multifarious activities in guideline formalization [19]. In contrast to that, much less attention has been devoted to a systematic comparison of guideline-based treatment processes with actual clinical treatment processes.

2.2 Analyzing Treatment Processes

The general strategy of comparing medical guidelines to clinical treatment practice is depicted in Fig. 2. First, a guideline is converted into a guideline-based process model composed of key actions. An actual treatment process depends on patient attributes X as well as institutional parameters θ reflecting the personal decisions and institutional environment of the acting health professional (cf. Subsection 2.1). With respect to the personal attributes we have to take into account that only a subset of these attributes is used as decision parameters for treatment according to the guidelines. Hence we split this patient attributes in diagnostic attributes X_d and personal attributes X_p, i.e. $X = (X_d, X_p)$. Individual attributes of a patient are denoted by (x_p, x_d). Treatment of a patient with diagnostic attributes x_p within a given institutional setting θ according to a guideline would result in so called *synthetic* log data denoted by $l_g(x_d, \theta)$. If we know the distribution $P(X_d)$ of the diagnostic attributes in a population of interest, we can simulate the distribution of possible synthetic treatment logs for this population for any given institutional setting θ. Accordingly, the random function of these synthetic logs is denoted by $l_g(X_d, \theta)$. In order to obtain the distribution $P(X_d)$, we use epidemiological data about the health status of the entire population, e.g., from prevalence studies.

Application of the guidelines in the treatment process of a patient depends usually not only on diagnostic attributes x_d but also on some of the personal attributes x_p. Hence, we denote in contrast to simulated process log data, clinical log data for a patient with attributes $x = (x_d, x_p)$ by $l_e(x_d, x_p, \theta)$, and the random function describing the empirical treatments for a patient collective by $l_e(X_d, X_p, \theta)$. Now, provided that both clinical and simulated synthetic log data are represented in a unified format (cf. Subsection 4.2), the problem of treatment compliance assessment can be (re-)stated as one of analyzing the deviance

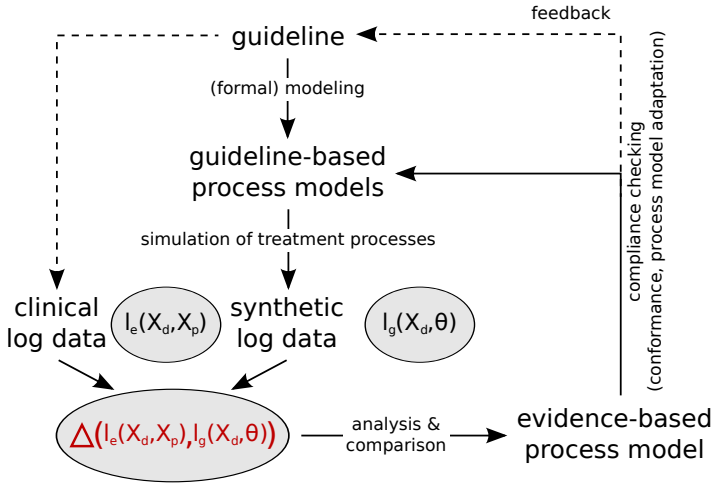

Fig. 2. Overall Methodology

$\Delta(l_e(X_d, X_p, \theta), l_g(X_d, \theta))$. More specifically, depending on available knowledge about $P(X_d)$ and the choice of the health system parameter, θ, typical research inquiries are as follows:

Distribution of clinical log data. Given X_d for a well-defined population, the degree of deviation of $l_e(X_d, X_p, \theta)$ from $l_g(X_d, \theta)$ can be analyzed, and whether such deviations can be explained in terms of X_p. Deviations of interest could be, e.g., partial compliance to the guidelines (such as patients modifying the intake of medication, delaying follow up, or exiting from after-care altogether). Formally, analysis – using traditional statistical techniques adapted to process data – results in distributions of clinical log data relative to specific patient sub-collectives defined by synthetic data.

Estimation and comparison of institutional parameters, θ, if unknown, from clinical log data. Prior to such an estimation, of course, each instance of clinical log data has to be assigned to one class of synthetic log data, using traditional machine learning methods in combination with process mining ([1, 18]).

Outcome analysis, seeking to figure out in how far realizations of different classes of synthetic log data, defined by the conditional process logic, influence the outcome of the treatment, and the sensitivity of the deviations from the clinical log data with respect to the outcome. Note that, contrary to traditional medical evidence (cf. Subsection 2.1), rather complex temporal treatment patterns (rather than well-defined treatments) have to be be evaluated using techniques of process mining, adding considerable analytical value compared to more conventional statistical approaches.

So far, we have formulated these research questions from a formal analytical point of view. Yet from practical point of view, evaluation of *medical conse-*

quences of these analyses is of utmost importance. In particular, the method may be used for evaluating which differences between synthetic treatment processes are of relevance from medical point of view. Eventually, as result of the analysis, *evidence-based process models* ensue, feeding back to the compliance evaluation of underlying medical guidelines.

Another issue of interest concerns the application of the approach for improving the quality of information about treatment processes. The simulated synthetic log data $l_g(X_d, \theta)$ based on a guideline give a rather complete picture of possible results. Comparison with the empirical treatment data $l_e(X_d, X_p, \theta)$ can identify incomplete medical treatment data as well as differences in the granularity of different data sources. In that way, the approach opens the opportunity to identify "blind spots" in documentation systems. Sometimes methods for data imputation and transformations for aligning different levels of detail can help to improve data quality for existing data about treatment processes.

2.3 Related Work

The EBMC^2 project in general and the methodology presented in this paper in particular are related to the areas of evidence-based medicine, medical guidelines, medical and healthcare processes, as well as process mining. Related work in the transition from evidence-based medicine to medical guidelines has been discussed in Section 2.1; medical and healthcare processes emerge as research topics in different domains. The Evimed project [12], for example, addresses literature research as to how guidelines relate to medical evidence. Apparently, these activities are rather orthogonal to the focus of the EBMC^2 project.

Process mining refers to a bundle of techniques to discover and analyze different facets of business processes. Some of these techniques have been applied to different application scenarios, particularly in the healthcare domain [17] where in case studies, mostly, hospital processes were discovered (mined) from process logs. Process discovery could be useful in the context of the EBMC^2 project as well, however, the main focus will be on *process synthesis* based on guidelines an delta analysis of synthetic and clinical processes. In general, simulation is an effective method to gain insights into real-world processes [1]. In the EBMC^2 project we investigate the applicability of simulation and delta analysis in the healthcare domain. Specifically, we aim at analyzing the causes for deviations between guideline-based processes and clinical processes.

3 Cutaneous Melanoma as a Case in Point

For the sake of illustration, a sample medical guideline concerning the diagnosis and treatment (but not prevention) of Cutaneous Melanoma [9] is used. Clearly, the narrative nature of the guideline addresses the physician in summarizing, based on studies and data available at the time of preparation, best practices of diagnostic approaches as well as a variety of established therapies including advanced stages of the disease. It is important to keep in mind that it "... may

be necessary or even desirable to deviate from these guidelines in the interest of specific patients or under special circumstances." ([9]: 271), emphasizing the generic character of the guideline as a set of – sometimes rather soft – constraints on advisable treatment decisions, and progressions, respectively. In spite of this generality, the overall structure of the guideline can be captured formally in the three-phase process skeleton exhibited in Fig. 1, of which Fig. 3 excerpts the aftercare phase expressed in a widespread standard of business process modeling, BPMN (Business Process Modeling Notation), chosen from amongst a set of equivalent alternatives such as GLIF, Asbru, EON, PROforma, GUIDE, or PRODIGY, all demonstrably "... close to traditional workflow languages" [16], as a direct comparison between BPMN and PROforma confirms that business process modeling languages can cope quite well with requirements of the health care domain [7].

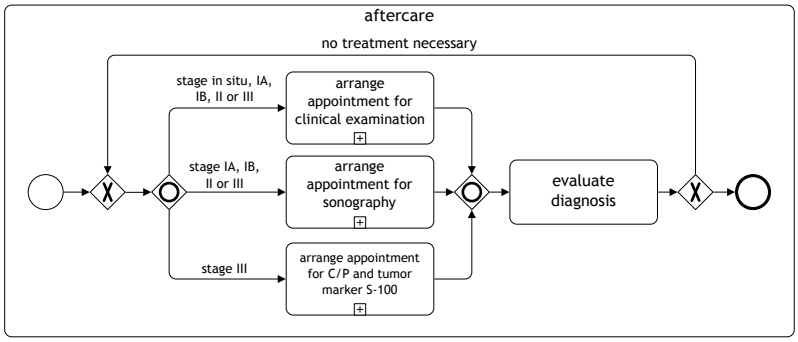

Fig. 3. BPMN Model: Melanoma Aftercare Phase

The example highlights both, the conditional branching and looping of progressions within a guideline as well as the nesting of subprocesses, admitting a successive refinement of process models while retaining their fundamental, non-recursive block structure. Fig. 4 illustrates refinement by nesting for the case of repeated *sonography* appointments in patient aftercare, stating the suggested choice of appointment intervals dependent on melanoma stage and the time lapse since last tumor diagnosis; again, however, it must be stressed that the guideline admits "considerable variation in follow-up approaches" ([9]: 279) because of an apparent lack of empirical data legitimating any particular recommendation, and that the depicted subprocess is but a debatable variant.

Regardless of the preferred modeling stance, particularly the branching decisions of treatment progressions may refer to patient characteristics and conditions beyond those stated explicitly in the guideline, like the staging of melanoma, tumor thickness, etc. ([9]: 273f). Accordingly, with respect to both, diagnostics and therapeutic actions, the representation language used to model patient

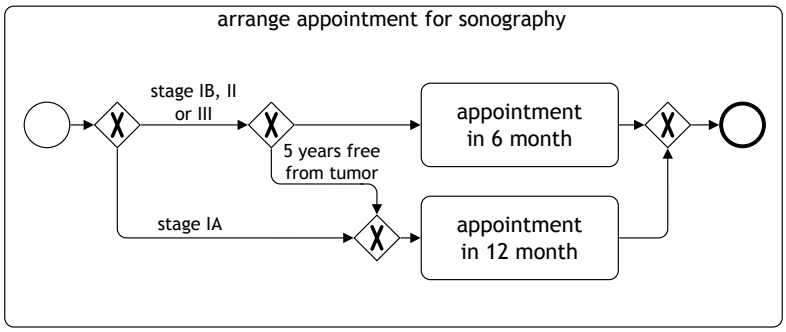

Fig. 4. BPMN Model: Sonography Aftercare Subprocess

treatment processes has to encompass a broader range of expressive elements, including typical clinical parameters (such as lab readings), patient characteristics (such as demographic variables, anamnesis, survival time, etc.), and medical interventions, in order to account appropriately for the cited "interest of specific patients or ... special circumstances." Moreover, if treatment process schemata are to be induced from medical patient records by means of process, or decision, mining methods, salient distinctions between actual treatment progressions, and branchings therein, may escape recognition simply because of an undue scarcity of formal expression.

4 Guideline-based Simulation and Analysis of Medical Treatment Processes

The overall methodology for modeling, simulating, and analyzing medical guidelines and clinical treatment processes is shown in Fig. 2. Deriving formal process models based on medical guidelines is illustrated by the Cutaneous Melanoma guideline in Section 3. In this section we focus on the generation of synthetic treatment processes based on the guideline process models and sketch how possible analysis and comparison of these data with clinical treatment processes can be conducted. Fig. 5 depicts the different levels: a) medical guidelines setting out a "frame" for possible treatment processes, b) synthetic treatment processes that reflect possible process executions based on the medical guidelines and can be created using simulation techniques, and c) the clinical processes, i.e., the processes implicitly executed or explicitly supported by a process-aware information system (PAIS) within the hospital or clinic. Comparing a) and c) leads to insights in how far the clinical treatment processes follow their corresponding guidelines and – in case they do not – find out about the reasons why. This corresponds to the questions set out by the area of business process compliance (BPC) [15]. Here focus is on developing (semi-)automatic techniques to check

compliance of real-world processes with guidelines, regulations, and compliance rules.

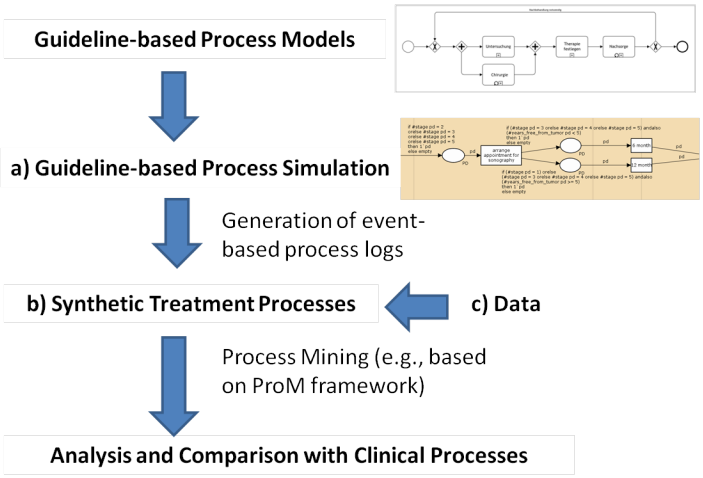

Fig. 5. Basic Modules of Simulation

Based on the guideline process models b) synthetic treatment processes can be generated using simulation tools such as CPN-Tools [13]. Generally, such simulation tools can take different parameters into account for which, e.g., certain probability distributions can be defined [2]. Hence, first of all, we can generate the possible interpretation of the guideline by generating possible process executions based on the guideline process models. As illustrated in Fig. 6B, the relevant process data elements are *stage* and *years_free_of_tumor*, since all decisions within the guideline depend on them. However, other clinical data might be available such as *age* or *gender*. Feeding these additional data into the simulation reveals possible influences on decision points within the treatment processes. Finally, comparing synthetic and clinical treatment processes b) and c) might yield additional insights if, e.g., synthesized processes considering additional data elements correspond, but also deviate from the real-world processes.

4.1 Generating Synthetic Evidence

As discussed before, for synthesizing treatment processes based on guideline process models, a suitable process simulation tool should offer facilities to define conditions to control the process flow depending on patient treatment data. Additionally, it is required that the tool is able to transport these patient treatment data through the process. Reviewing possible alternative tools for process simulation [13] and the aim for using the synthesized process data for process

mining we decided to use CPN-Tools for our illustration. We use the approach described by [2] to create computable log files out of the CPN-Tools simulation.

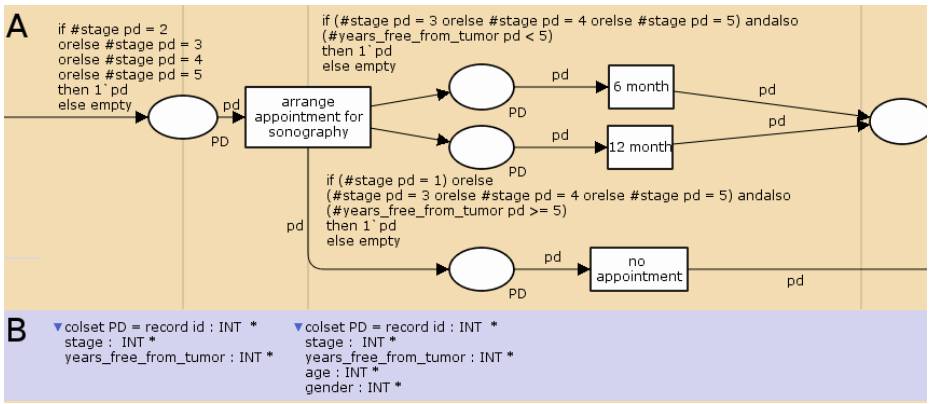

Fig. 6. Colored Petri net model in CPN-Tool: Sonography Appointment Subprocess

Figure 6A shows the subprocess (cf. Fig. 4) now modeled as a colored Petri net within CPN-Tools. We first use a record-typed color set (cf. Fig. 6B) as a token that holds patient data which include the two parameters *stage* and *years_free_of_tumor* that are used for the decisions within the Melanoma guideline. A simulation on this Petri net followed by analyzing with process mining tools will result in a guideline compliant process model.

As motivated in Section 3, typically a patient data record includes more information than considered within the guideline. To simulate this fact we add additional data elements such as *age* and *gender* to our simulation by extending the patient data record (cf. Fig. 6B). As a result we can determine if the process mined model will differ from the guideline, moreover we can identify where and why the simulated – and in case of real-world data the real – treatment processes deviate from the guideline. Note that another possible way to simulate processes that deviate from the guideline is an alteration of the colored Petri net by manipulating the model itself (e.g., by adding a 9 month activity for arranging an appointment). Finally, the simulation model can be extended by "predictable" exceptions that are not included within the guideline, but are valuable to analyze, e.g., patient behavior. In Melanoma aftercare, one possibility is that no arrangement for the next check-up is made since the patients may not show up (see the additional exit transition "no appointment" in Fig. 6).

4.2 Representing Process Evidence

The comparison and statistical analysis of both, clinical and synthetic treatment process data, necessitates a simple yet versatile data structure using some unified,

flexible format capable to represent all process-contingent patient data. Essentially, a *medical biographic store* (MBS, for short) is arranged as a logical data frame recording process log data originating from both, clinical care documentation (empirical log data from clinical treatment processes, such as electronic patient records) and simulation runs using guideline-based process models. More specifically, with respect to the notation introduced above, a clinical treatment process produces, for some patient taken care of, a trace $y(x_d, x_p)$ of log data, whereas "applying" some guideline-based process model g to this patient yields a trace of synthetic log data $g(x_d, \theta)$.

The MBS relates to electronic health record approaches (cf. ISO/CEN EN 13606; [4, 6, 8, 11, 14]), yet without committing itself to any particular syntactic standard of health data representation. Rather, the data structure of the MBS annotates each elementary entry, or *protocol particle*, with a spatio-temporal and subject-matter reference to account for a variety of modes of subsequent data selection and aggregation operations, yielding a quadruple structure

$$\langle pid, \langle temp, loc, type\rangle, feat, value\rangle$$

where *pid* signifies an (MBS-)internal patient identifier, *temp* the temporal and *loc* the location references, respectively, of the *value* field, while *type* annotates the origin, or context (such as diagnosis, medication, intervention, ...) of the recorded *value*, and *feat* represents the clinical, or medical, feature (of the patient *pid*) the *value* is the recorded value of. Ontologically, protocol particles record *partial* discontinuities of patient (health) states, assuming that absence of such evidence implies the steadiness of state. As an example from the application domain, a protocol particle

$$\langle \#2006.184.277, \langle 2007\text{-}09\text{-}17, AKH\ Wien\text{-}Derma, dgn\rangle, melanoma\text{-}stage, IIC\rangle$$

would record the clinical documentation entry for (real) patient *#2006.184.277* diagnosing a stage *IIC* melanoma on *Sept. 17, 2007*, at the *Dermatological Clinic* of the General Hospital Vienna. Technically, each particle component may in turn bear a rich (that is, nested) formal substructure as mandated by particular data-analytical investigations, notably process and decision mining to discover actual patient treatment process schemata, and entailed by the semantics of *feat* entries.

For purposes of statistical analysis, it is mandatory that case-by-variate structures can be derived easily from the MBS by (i) converting MBS extractions into "flat file" views by simply selecting subsets of *feat* fields recorded in protocol particles, and (ii) facilitating the formal addition of further variables by defining various temporal etc. predicates on the annotation fields of a patient's protocol particles (e.g., in order to investigate the temporal dynamics of an event sequence, etc.).

4.3 The Analysis Dimension: Compliance (analytical techniques)

After simulating the guideline with $l_g(X_d, \theta)$ the first hypothesis is that there are no differences in the clinical log data with respect to personal attributes

such as *age* and *gender* (cf. Fig. 6). One technique to analyze this hypotheses on decision points and possible exit states within the synthetic treatment processes is provided by decision mining [24]. Decision mining is implemented within the process mining framework ProM [16]. Using the reporting tools of ProM we can produce the following report table about the outcomes of the treatment processes (cf. Table 1), e.g., age and gender groups. Under the hypotheses there should be no significant differences in the exit state distributions, i.e., dead, healed, no show, keep appointment. Otherwise one has to think about reformulation of the guidelines in terms of personal attributes X_p or improving the implementation strategy of the treatment process.

Table 1. Example Report Structure

	dead	healed	no show	keep appointment
age < 20	x %	...	...	...
$20 \leq$ age ≤ 30	y %	...	...	...
...	...	...	...	...
gender = f	z %	...	...	...

5 Conclusions and Outlook

The virtues of the proposed approach towards assessing the guideline-compliance of clinical treatment processes can be summarized as follows: (i) formalizing the representation of clinical treatment processes amounts to a rigorous process view on medical evidence; (ii) formalizing clinical treatment processes as well as medical guidelines allows to restate the concept of process compliance in a more precise fashion than before; and (iii) evidence-based data-driven compliance assessment highlights the actual scope of guideline implementations, whether compliant or not. Obviously, formal compliance assessment drives the understanding of clinical practice, and, in doing so, feeds back to the maintenance and evolution of guidelines in the light of empirical evidence.

As it stands, though, the outlined methodology is but a research program, entailing a range of research issues including, to name just the most important ones, the development of a modeling framework for patient treatment processes accounting for a wide variety of conditions and exceptions typically not even mentioned in the narratives of medical guidelines, effective approaches towards wrapping and integrating available patient record data from various – and often scattered – clinical documentation and electronic health record infrastructures, strategies for deciding about the scope of medical evidence to be included in the analyses at all, the tuning of the simulation apparatus generating synthetic evidence from formal process models for selected patient population collectives, and both, statistical process mining and decision mining techniques, applied to available data bodies.

Obviously, theoretical advancement of the topic will have to be paralleled by extensive practical experimenting with a range of medical guidelines in the light of different bodies of medical and epidemiological evidence. For the time being, within the EBMC^2 project, three valuable data sources are investigated thoroughly, namely (i) a detailed data collection of clinical Cutaneuos Melanoma stage 4 protocols, (ii) a vast body of administrative data of the Austrian Main Association of Social Insurers comprising a comprehensive picture of medical patient treatments, and (iii) Melanoma-relevant excerpts of the Austrian cancer register. Currently, much effort is devoted to the formal process-focused integration of these (as well as further) sources of patient-related Melanoma data as an indispensable, yet non-trivial prerequisite to process mining and, at a more general level, to the validation of the proposed EBMC^2 analytical methodology of comparing medical treatment practice with medical guidelines.

Acknowledgements. We thank all our collaborators in the EMBC^2 project: Michael Binder, Wolfgang Dorda, Georg Duftschmid, Walter Gall, Milan Hronsky, Christoph Rinner, and Stefanie Weber.

References

1. van der Aalst, W.M.P.: Process Mining – Discovery, Conformance and Enhancement of Business Processes. Springer, New York-Berlin (2011)
2. Alves De Medeiros, A.K., Günther, C.W.: Process mining: Using CPN tools to create test logs for mining algorithms. In: Proceedings of the Sixth Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools. pp. 177–190 (2005)
3. Benthem, J.V.: The Logic of Time. Reidel, Dordrecht (1983)
4. Blobel, B.: Advanced EHR architectures – promises or reality? Methods Inf Med 45(1), 95–101 (2006)
5. Cochrane, A.: Effectiveness and Efficiency – Random Reflections on Health Services. Royal Society of Medicine Press, London, UK (1972)
6. Eichelberg, M., Aden, T., Riesmeier, J.: A survey and analysis of electronic healthcare record standards. ACM Computing Surveys 37(4), 277–315 (2005)
7. Fox, J., Black, E., Chronakis, I., Dunlop, R., Petkar, V., South, M., Thomson, R.: From guidelines to careflows: Modelling and supporting complex clinical processes. In: Computer-based medical guidelines and protocols: a primer and current trends, pp. 44–61. IOS Press, Netherlands (2008)
8. Gall, W., Grossmann, W., Duftschmid, G., Wrba, T., Dorda, W.: Analyses of EHRs for research, quality management and health politics. Proc. MIE2008 XXIst International Congress of the European Federation for Medical Informatics / Studies in Health Technology and Informatics 136, 425–430 (2008)
9. Garbe, C., Peris, K., Hauschild, A., Saiag, P., Middleton, M., Spatz, A., Grob, J., Malvehy, J., Newton-Bishop, J., Stratigos, A., Pehamberger, H., Eggermont, A.: Diagnosis and treatment of melanoma: European consensus-based interdisciplinary guideline. European Journal of Cancer 46(2), 270–283 (2010)
10. Guyatt, G., Oxman, A., Schüneman, H., Tugwell, P., Knottnerus, A.: Grade guidelines: A new series of articles in the journal of clinical epidemiology. Journal of Clinical Epidemiology 64, 380 – 382 (2011)

11. Häyrinen, K., Saranto, K., Nykänen, P.: Definition, structure, content, use and impacts of electronic health records: A review of the research literature. Int J Medical Informatics 77, 291–304 (2008)
12. Hunter, A., Williams, M.: Using clinical preferences in argumentation about evidence from clinical trials. In: Proc. 1st ACM Int'l Health Informatics Symposium (2010)
13. Jansen-Vullers, M., Netjes, M.: Business process simulation – a tool survey. In: Workshop and Tutorial on Practical Use of Coloured Petri Nets and the CPN Tools, Aarhus, Denmark (2006)
14. Koch, O.: Process-based and context-sensitive information supply in medical care. In: Proc. 2nd Int. Workshop on Context-Based Information Retrieval. Computer Science Research Report 114, pp. 93–104. Roskilde University, Denmark (2007)
15. Ly, T., Rinderle-Ma, S., Göser, K., Dadam, P.: On enabling integrated process compliance with semantic constraints in process management systems – requirements, challenges, solutions. Information Systems Frontiers pp. 1–25 (2009)
16. Mans, R.: Workflow Support for the Healthcare Domain. Proefschriftmaken.nl, Netherlands (2011)
17. Mans, R.S., Schonenberg, M.H., Song, M., Aalst, W.M.P., Bakker, P.J.M.: Application of process mining in healthcare – a case study in a dutch hospital 25, 425–438 (2009)
18. Marsland, S.: Machine Learning – An Algorithmic Perspective. Chapman & Hall/CRC , Boca Raton (2009)
19. OpenClinical: http://www.openclinical.org/gmmintro.html, accessed 28-07-2011
20. Oxford Centre for Evidence-Based Medicine: http://www.cebm.net/index.aspx?o=5653, accessed 28-07-2011
21. Perleth, M., Busse, R., Gerhardus, A., Gibis, B., Luhmann, D.: Health Technology Assessment: Konzepte, Methoden, Praxis für Wissenschaft und Entscheidungsfindung. Berliner Schriftenreihe Gesundheitswissenschaften, Medizinisch Wissenschaftliche Verlagsgesellschaft (2008)
22. Rescher, J., Urquhart, A.: Temporal Logic. Springer, New York-Berlin (1971)
23. Rosenfeld, R., N.Shiffman, R.: Clinical practice guideline development manual: A quality driven approach for translating evidence into action. Otolaryngol Head Neck Surg. 140(6 Suppl 1, 1– 43 (2009)
24. Rozinat, A., van der Aalst, W.M.P.: Decision mining in ProM. In: Business Process Management. pp. 420–425 (2006)
25. Shiffman, R.N., Karras, B., Agrawal, A., Chen, R., Marenco, L., Nath, S.: Gem: a proposal for a more comprehensive guideline document model using xml. Journal of the American Medical Informatics Assoc. 7, 488 – 498 (2000)
26. The Cochrane Collaboration: http://www.cochrane.org/, accessed 28-07-2011