Ontological Modelling of a Psychiatric Clinical Practice Guideline

Daniel Gorín[1], Malte Meyn[1], Alexander Naumann[2], Miriam Polzer[1], Ulrich Rabenstein[1], and Lutz Schröder[1]

[1] Friedrich-Alexander-Universität Erlangen-Nürnberg
[2] Psychiatrische Klinik Lüneburg

Abstract. Clinical practice guidelines (CPGs) serve to transfer results from evidence-based medicine into clinical practice. There is growing interest in clinical decision support systems (CDSS) implementing the guideline recommendations; research on such systems typically considers combinations of workflow languages with knowledge representation formalisms. Here, we report on experience with an OWL-based proof-of-concept implementation of parts of the German S3 guideline for schizophrenia. From the information-technological point of view, the salient feature of our implementation is that it represents the CPG entirely as a logic-based ontology, without resorting, e.g., to rule-based action formalisms or hard-wired workflows to capture clinical pathways. Our current goal is to establish that such an implementation is feasible; long-range benefits we expect from the approach are modularity of CPG implementation, ease of maintenance, and logical unity.

1 Introduction

Clinical practice guidelines (CPGs) are consensus documents intended to improve the quality of the treatment of specific diseases following the paradigm of evidence-based medicine. In the current work, we develop a methodology for the ontology-centered implementation of CPGs that represents the guideline content uniformly in terms of a description-logic based ontology, rather than having separate expressive means for actions, e.g. a rule base as in [9] or a dedicated workflow formalism as in many guideline languages (Section 5). We present a proof-of-concept stage CPG-based clinical decision support system (CDSS) for schizophrenia implemented following this methodology, formalizing parts of the German S3 (i.e. high-evidence) guideline for schizophrenia [5]. Envisaging our framework to evolve into a generic Clinical Guideline Module, currently instantiated to ICD10 diagnostic code F20 (Schizophrenia), we call the arising tool CGM/F20. Long-term advantages we expect from a purely ontological approach include logical coherence, higher flexibility than achievable by hard-wired workflows, and ease of maintenance. We do emphasize that our preliminary results do not yet justify any claim that these benefits will actually be realized; the purpose of the current work is to explore how far a purely declarative logic-based implementation of a CPG is possible in principle. The schizophrenia guideline is particularly suitable as a case study, being given largely as a collection of recommendations on good clinical practice rather than as a set of definite clinical algorithms.

2 An Assistant for the Diagnosis and Treatment of Schizophrenia

We briefly describe the envisaged use of CGM/F20[3]. The tool consists of two parts, a generic clinical guideline module, CGM, and a formalization of (portions of) the *Schizophrenia Guideline* of the German Association of Psychiatry, Psychotherapy and Psychosomatics (DGPPN) [5]. The tool is not currently integrated into a real HIS, and instead implements basic patient data management itself in a fairly naive way to enable the intended proof of concept.

Recommendations are based on the medical record of a patient, represented as a time-ordered sequence of events such as symptoms reported, diagnoses made, test results, therapeutic measures etc. To see how this works, let us assume that a patient has been recently admitted and it is clear that he has been hearing hallucinatory voices commenting on his behaviour for a prolonged time. Because of additional symptomatology and patient background, he is considered possibly HIV positive. The DGPPN-SG stipulates [5, pp. 31–32]:

> *A diagnosis of schizophrenia requires at least one* unmistakable *symptom from Groups 1 to 4 (or two, if they are less certain), or at least two symptoms from Groups 5 to 8. These symptoms have to be constantly present for a month or longer.*

(Translated from the original German, like all further guideline quotes; our emphasis.) So let us assume that the patient's EHR already contains the above-mentioned information that he "unmistakably" hears commenting voices. Since according to the DGPPN-SG, this is a symptom in Group 3, the system should recommend the diagnosis of schizophrenia. In fact, the system mentions "acute schizophrenia" as a possible diagnosis, but indicates that the recommendation cannot yet be made definite due to missing information. This follows from the guideline, which not only lists exceptional cases under which the diagnosis should not be made [5, p. 32]:

> *One should not diagnose schizophrenia in the presence of an unmistakable brain disorder or while the patient is intoxicated or undergoing detoxification*

but also dictates a differential diagnosis to rule out alternative psychotic disorders. In addition to the putative diagnosis, the system recommends a *measure*, namely to perform a test for HIV; this again follows from the guideline:

> *In case of a corresponding suspicion, an HIV test [. . .] should be performed.*

At this point the physician could, e.g., immediately dismiss "detoxification" as a possibility, and then order blood tests and a brain scan to rule out the remaining obstacles to the diagnosis. If the patient's medical record were then updated with negative results for all tests, the other scenarios would vanish and the system would then actually recommend the diagnosis of acute schizophrenia.

We stress that this would be just a recommendation, that is, the diagnosis would not be automatically issued by the system. So assume the physician agrees with this diagnosis and records it in the system. The DGPPN-SG will now recommend a pharmacotherapy, but only once the informed consent of the the patient is recorded, as specified in the guideline [5, p. 43].

[3] CGM/F20 is open source, and available at `http://www8.cs.fau.de/research/cgm`.

3 Implementation Challenges

We discuss challenges met in the implementation of the guideline (with solutions described in Section 4), and pinpoint aspects that make an ontology-based approach fit particularly well, but also ones were it becomes apparent that the current state of the art in description logics could be improved.

Partial knowledge. Unlike database systems, the system should follow the *open world assumption* as used in logic-based ontology languages; e.g. a condition being unmentioned in a patient's health record should be not conflated with absence of the condition.

Defaults and exceptions. In Section 2 we saw a guideline excerpt formulated in *default case/exceptional case* style: The CPG recommends to diagnose schizophrenia when sufficiently many symptoms from a given list are present, but then lists *exceptions* to this rule, such as not to issue the diagnosis if the patient has a brain disorder. Under the open world assumption, just incorporating absence of exceptions as a condition into the general rule would not yield the expected results: For instance, as we cannot yet dismiss the possibility of a brain disorder at the time the patient with commenting voices in Section 2 is admitted, the system would then simply not recommend a diagnosis of schizophrenia, without further comment. It is thus crucial to distinguish between default actions and exceptions when formalizing a guideline. On the other hand, the intended behaviour is not quite the same as in standard default reasoning paradigms, which in the example would presumably simply recommend diagnosing schizophrenia without further comment, unless a brain disorder is *known* to be present. In contrast, our system would indicate to the physician that schizophrenia is a likely diagnosis but explicitly hold off on a recommendation until exceptional conditions are known to be absent.

Epistemic queries. The user interaction with the system should depend on whether certain facts regarding a patient are *known* or not; e.g. the system should not ask questions whose answer is already in the patient record, or pursue possible diagnoses that are already known to be ruled out. A full treatment of such situations requires a query language with an epistemic operator [11], not currently supported by mature OWL reasoners. We deal with this issue by letting the tool launch appropriate ontological queries to check whether certain facts are known, and then recombine queries, in particular *negatively*.

Dismissing possible exceptions. Exceptions may rest on complex logical combinations of facts (e.g., a medication can be counterindicated by a combination of conditions); dismissing such an exception as a whole amounts to negating a complex concept, so having an expressive logic with unrestricted negation helps.

Judicious information gathering. Recommendations of the schizophrenia guideline sometimes depend on symptoms being "unmistakable" (*eindeutig*), and the system needs to query the physician accordingly. On the other hand, the system should not issue such a query for other symptoms (such as 'high fever').

Temporal relations The guideline needs information about the temporal order and duration of certain events, and moreover has concepts of a *current phase* and *previous phases* of the disease.

4 The CGM/F20 Ontology

The CGM/F20 ontology consists of two static OWL ontologies, one for a general modelling of time and one for the actual guideline content (once additional CPGs are covered in future extensions, the latter will split into a generic and a disease-specific part). Additionally, each patient is modelled as a separate dynamic ontology that extends the static ontologies with assertional knowledge modelling the patient data; it is generated at admission and is continuously updated by the system as the clinical process progresses.

Time As the guideline only needs fairly large-scale temporal concepts, we opt for modelling a highly abstracted temporal ontology from scratch (partly using SWRL) rather than import more fine-granular (clinical) temporal ontologies such as CNTRO [18]. Our time ontology is based on the central concept of *event*. Events may have a beginning, an end, and a duration, and may belong to a *session*. We use sessions as logical time units, ending with the release of the patient. One phenomenon creating the need for sessions is that many guideline statements (e.g. on specific laboratory tests and on efficiency and side effects of medications) depend on schizophrenia manifesting itself *for the first time*. An attempt to model this condition by saying that the patient was not diagnosed with schizophrenia in the past will clearly fail, as the relevant guideline recommendations will then be blocked from the moment the patient is diagnosed with schizophrenia. Our solution is to formulate the condition in terms of sessions – we require that the patient has not been diagnosed with schizophrenia in any previous session.

Overall Design of the Guideline Model Unsurprisingly, the guideline model revolves around the concepts of Symptom, Diagnosis, Measure, and Patient . Additionally, there is a class TherapyGoal modelling abstract overall goals of the therapy that are not refined to concrete therapeutic measures in the guideline. The recommendations of the guideline are implemented in terms of named classes of patients (those to which the recommendation applies), the *recommendation classes*. These concepts are linked with the time ontology; in particular, symptoms, diagnoses and measures are events in the sense of the time ontology, and always belong to a session.

One important point to think about in designing an ontology is what individuals one imagines as potential members of a given class. For instance, we regard classes of symptoms and measures as inhabited by concrete instances attached to a given person over a given time period, e.g. 'the delusions experienced by patient X in the past three weeks'. To avoid conceptual pitfalls as pointed out in [15], not all inhabitants of such classes are required to exist in the real world – e.g. a therapeutic measure might just be recommended or currently planned. How far such entities are real or just putative is determined by their relationship to other individuals.

Recommendations are implemented in terms of object properties linking measures with patients, and governed by definitions of recommendation classes. There are three types of such properties: recommendations, non-recommendations (e.g. for diagnoses already issued in the past), and counter-recommendations (explicit discouragements). As indicated in Section 3, we need to emulate a form of default reasoning. To this end, recommendation classes may, via annotations, be associated with preconditions,

expressed by *precondition classes*. This enables us to assign patients to a recommendation class although we yet do not know whether they fulfill these additional preconditions, and then confirm the recommendation once the preconditions are checked.

Annotations are used to configure the behaviour of the tool, in particular the query mechanism and the user interface. Besides for preconditions, we use annotations to trigger queries for additional qualifications of symptoms such as unmistakability; to link to the guideline text; and to mark preconditions that the physician is allowed to summarily dismiss without justification.

Reasoning The reasoning service by means of which CGM/F20 arrives at recommendations is querying via SPARQL-DL formulas [17]. Ideally, one would want to query for classes C of diagnoses such that the patient at hand belongs to classes such as *hasRecommendedDiagnosis* **some** C; however, constructs of this type go beyond (the well-documented and stable part of) SPARQL-DL. As a workaround, we populate classes of measures and diagnoses with generic individuals, which we use in place of the class in the definition of recommendation classes. We then query for individuals that are, say, diagnoses, and are related to the patient at hand via *hasRecommendedDiagnosis*. For example, diagnostic recommendations for patient *John* are produced by the query

SELECT ?diagnosis **WHERE** {
John schizophrenia:*hasRecommendedDiagnosis* ?diagnosis.
FILTER NOT EXISTS {John schizophrenia:*hasNonRecommendedDiagnosis* ?diagnosis}.}

That is, we first query for recommended diagnoses and then filter out the non-recommended ones in our emulation of defeasible reasoning. In the example, if John has an unmistakable symptom from a certain list of symptoms, he will satisfy *hasRecommendedDiagnosis* R_Schizophrenia, but the diagnosis will not be recommended if that diagnosis has already been made, putting John also in the class *hasNonRecommendedDiagnosis* R_Schizophrenia, so that R_Schizophrenia is removed from the result of the query. Having thus established diagnoses to be explored, we next need to check the status of the preconditions before a diagnosis can be conclusively recommended. We thus query for all classes John belongs to, and then iterate over all preconditions of recommendation classes John belongs to and check their status: If the John belongs to the precondition class, its status is 'confirmed'; if he is in the negation of the precondition class, the status is 'excluded'; and otherwise 'unknown'.

Modularity Currently we are only modelling a single CPG, but we provide an interface which should also fit for other guidelines. In a multi-guideline framework, one would put general-interest classes and object properties revolving around symptoms, diagnoses, therapy goals etc. into a master ontology and let ontologies for specific CPGs build on this ontology, using the same types of annotations. The tool would likely require only little adaptation, barring extensions necessary to cover new general phenomena not occurring in the schizophrenia guideline. In a multi-guideline framework, guidelines could and should refer to each other, e.g. for differential diagnoses. The logic-centered approach would then presumably play out its advantages quite visibly, as clinical pathways could be generated by combining logical axioms across different guidelines, avoiding the need for laborious manual integration of workflows.

5 Related Work

Our work is situated in the highly active area of clinical decision support. For systematic reviews of clinical decision support systems (CDSS) and guideline implementations, see [7, 12]. Current CDSS are mostly focused on somatic diseases, one exception being the CompTMAP tool that implements a hard-wired medication protocol for depression [19]. A range of dedicated formalisms has been developed for the implementation of clinical practice guidelines; see [13] for an overview. A common denominator of these languages is that they focus on workflow descriptions fixing specific clinical paths, and use knowledge representation only at hardwired decision points in workflows. More recently, CDSS have emerged that are, like CGM/F20, based primarily on ontological representations, typically in OWL. One approach is to represent workflows explicitly by modelling states and transitions in OWL and SWRL [20, 8, 6]. Beyond this, ontologies have been used for patient classification in diabetes [9] and lung cancer [16], and for decision support in breast cancer follow-up [1], comorbidities in cardiology [2], pre-operative testing [4], and chemotherapy [3]. Out of these, it is probably [4] whose approach is most closely related to ours. The main differences between the CDSS described in [4] and CGM/F20 are on the one hand the fact that CGM/F20 deals with extended clinical processes and hence needs to consider temporal concepts, and on the other hand specific implementation challenges addressed in CGM/F20 that lead to extended use of DL querying in CGM/F20, cf. Sections 3 and 4.

6 Conclusions and Future Work

We have explored an approach to CPG implementation that relies on representing the full content of the guideline, including procedural parts, as an OWL ontology. Our system CGM/F20 currently covers part of the German S3 guideline for schizophrenia. In its present state of development, CGM/F20 delivers recommendations for the next step in a clinical path (of course still leaving the actual decisions on diagnostic and therapeutic steps to the psychiatrist). The guideline content and, in particular, its appropriate interactive presentation to the user poses a number of specific challenges that are addressed in particular by the careful design of DL queries and intensive use of annotations.

As an immediate next step, still within the proof-of-concept stage, we will extend the tool to enable explicit look-ahead, i.e. to generate multi-step clinical pathways and replan these as the actual clinical process unfolds, possibly using OWL-S [10]. Moreover, we intend to substantiate our claim of genericity of our CPG implementation framework by formalizing additional guidelines, aiming especially for diseases whose diagnosis and treatment involve a vocabulary of concepts that intersects with those of schizophrenia. A point of particular interest is in-depth coverage of differential diagnoses. Interdisciplinary integration of guidelines will increase the need for connecting CPG ontologies with standard medical terminologies such as SNOMED CT or OpenGALEN [14].

References

1. S. Abidi, S. Abidi, S. Hussain, and M. Shepherd. Ontology-based modeling of clinical practice guidelines: A clinical decision support system for breast cancer follow-up interventions at primary care settings. In *MEDINFO 2007*, pp. 845–849. IOS Press, 2007.

2. S. Abidi, J. Cox, S. Abidi, and M. Shepherd. Using OWL ontologies for clinical guidelines based comorbid decision support. In *Hawaii International International Conference on Systems Science, HICSS 2012*, pp. 3030–3038. IEEE Comp. Soc., 2012.
3. C. Beierle, L. Eisele, G. Kern-Isberner, R. Meyer, and M. Nietzke. Using ontological knowledge about active pharmaceutical ingredients for a decision support system in medical cancer therapy. In *KI 2016*, vol. 9904 of *LNCS*, pp. 119–125. Springer, 2016.
4. M.-M. Bouamrane, A. Rector, and M. Hurrell. A hybrid architecture for a preoperative decision support system using a rule engine and a reasoner on a clinical ontology. In *Web Reasoning and Rule Systems, RR 2009*, vol. 5837 of *LNCS*, pp. 242–253. Springer, 2009.
5. Deutsche Gesellschaft für Psychiatrie, Psychotherapie und Nervenheilkunde. *Behandlungsleitlinie Schizophrenie*. Springer, 2006.
6. C. Doulaverakis, V. Koutkias, G. Antoniou, and I. Kompatsiaris. Applying SPARQL-based inference and ontologies for modelling and execution of clinical practice guidelines: A case study on hypertension management. In *Knowledge Representation for Health Care, KR4HC/ProHealth 2016*, vol. 10096 of *LNCS*, pp. 90–107. Springer, 2017.
7. A. Garg, N. Adhikari, H. McDonald, M. Rosas-Arellano, P. Devereaux, J. Beyene, J. Sam, and R. Haynes. Effects of computerized clinical decision support systems on practitioner performance and patient outcomes: a systematic review. *JAMA*, 293:1223–38, 2005.
8. B. Jafarpour, S. R. Abidi, and S. S. R. Abidi. Exploiting semantic web technologies to develop OWL-based clinical practice guideline execution engines. *IEEE J. Biomed. Health Inform.*, 20:388–398, 2016.
9. V. Kashyap, A. Morales, and T. Hongsermeier. On implementing clinical decision support: achieving scalability and maintainability by combining business rules and ontologies. In *AMIA Annual Symposium*, vol. 2006, pp. 414–418. AMIA, 2006.
10. D. Martin, M. Burstein, D. McDermott, S. McIlraith, M. Paolucci, K. Sycara, D. McGuinness, E. Sirin, and N. Srinivasan. Bringing semantics to web services with OWL-S. In *World Wide Web, WWW 2007*, pp. 243–277, 2007.
11. A. Mehdi, S. Rudolph, and S. Grimm. Epistemic querying of OWL knowledge bases. In *Extended Semantic Web Conf., ESWC 2011*, vol. 6643 of *LNCS*, pp. 397–409. Springer, 2011.
12. M. Peleg. Computer-interpretable clinical guidelines: A methodological review. *Journal of Biomedical Informatics*, 46:744–763, 2013.
13. M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R. A. Greenes, R. Hall, P. Johnson, N. Jones, A. Kumar, S. Miksch, S. Quaglini, A. Seyfang, E. Shortliffe, and M. Stefanelli. Comparing computer-interpretable guideline models: a case-study approach. *J. AMIA*, 10:52–68, 2003.
14. A. Rector, J. Rogers, P. Zanstra, and E. van der Haring. OpenGALEN: Open source medical terminology and tools. In *AMIA Annual Symposium*, p. 982. AMIA, 2003.
15. S. Schulz, H. Stenzhorn, M. Boeker, , and B. Smith. Strengths and limitations of formal ontologies in the biomedical domain. *Rev. Electron. Comun. Inf. Inov. Saude*, 3:31–45, 2009.
16. M. Sesen, R. Banares-Alcántara, J. Fox, T. Kadir, and J. Brady. Lung Cancer Assistant: An ontology-driven, online decision support prototype for lung cancer treatment selection. In *OWL: Experiences and Directions Workshop, OWLED 2012*, 2012.
17. E. Sirin and B. Parsia. SPARQL-DL: SPARQL query for OWL-DL. In *OWL: Experiences and Directions, OWLED 2007*, vol. 258 of *CEUR Workshop Proc.*, 2007.
18. C. Tao, W. Wei, H. Solbrig, G. Savova, and C. Chute. CNTRO: A semantic web ontology for temporal relation inferencing in clinical narratives. In *Ann. Symp. AMIA*, pp. 787–791, 2010.
19. M. Trivedi, J. Kern, B. Grannemann, K. Altshuler, and P. Sunderajan. A computerized clinical decision support system as a means of implementing depression guidelines. *Psych. Services*, 55:879–885, 2004.
20. Y. Ye, Z. Jiang, X. Diao, D. Yang, and G. Du. An ontology-based hierarchical semantic modeling approach to clinical pathway workflows. *Comp. Bio. Med.*, 39:722–732, 2009.