Reasoning with Effects of Clinical Guideline Actions Using OWL: AL Amyloidosis as a Case Study

Mor Peleg[1], Samson W. Tu[2], Giorgio Leonardi[3], Silvana Quaglini[3], Paola Russo[4], Giovanni Palladini[4], and Giampaolo Merlini[4]

[1] University of Haifa, Haifa, 31905, Israel
morpeleg@is.haifa.ac.il
[2] Stanford University, Stanford, CA, 94305, USA
swt@stanford.edu
[3] University of Pavia, Pavia, Italy
giorgio.leonardi@unipv.it, silvana.quaglini@gmail.com
[4] Amyloidosis Research and Treatment Center and Dept. of Biochemistry, IRCCS Policlinico San Matteo FDN and University of Pavia, Pavia, Italy
paola.rnm.russo@gmail.com,
{giovanni.palladini,gmerlini}@unipv.it

Abstract. We developed an ontology that allows representation and reasoning with effects of clinical actions. The ontology can support three important use-cases: (1) summarization and explanation of observed clinical states, (2) enhancing patient safety using safety rules, and (3) assessing guideline compliance. In this paper we focus on explanation of observed clinical states based on abductive reasoning that utilizes a causal network. We demonstrate our approach using examples taken from a guideline for management of amyloidosis.

Keywords: OWL, ontology, computer-interpretable guidelines, causal models.

1 Introduction

Motivated by the need to develop a guideline formalism that would allow easy maintenance of changing clinical knowledge and justification and explanation of recommendations, we developed an ontology that represents and reasons with abstractions. Abstractions, or high-level principles used by the guideline authors do not change rapidly, thus working at their level could potentially allow changing guideline details (such as adding a particular new chemotherapeutic drug to a guideline that already uses chemotherapy or changing the dose of an existing drug) without significantly changing the overall structure of the guideline. Moreover, using abstractions such as clinical states, actions, and relationships such as effects of clinical actions, allows us to represent mechanisms of actions of different therapies and the relationships between patient data that should be collected and particular diagnostic tests that could be used to collect them. These mechanisms of actions and relationships' are likely to change less during guideline updates [1]; what will mostly

D. Riaño, A. ten Teije, and S. Miksch (Eds.): KR4HC 2011, LNAI 6924, pp. 65–79, 2011.

change would be additions or removal of particular drugs belonging to established drug groups (whose mechanisms of actions were already modeled), dose changes, additions of laboratory tests, and criteria for selecting existing actions. These updates are necessary whenever results from new clinical trials suggest new treatments, for example a new drug is discovered and proved to be more effective than the existing ones, or whenever a new, more accurate, diagnostic procedure is assessed.

Representing abstractions, mechanisms of actions, and causal relationships allows us to perform **qualitative** reasoning to provide explanations and justifications for the guideline's recommendations thus ensuring that the guideline contains a complete and appropriate set of recommendations.

We exemplify our approach using examples taken from a therapeutic guideline and an electronic medical record (EMR) developed by the Italian Amyloidosis Study Group (SIA) for systemic immunoglobulin light chain amyloidosis (AL) [2]—a rare disease whose clinical knowledge needs to rapidly be updated according to the results of the ongoing clinical trials by the consensus panels of the dedicated medical societies. The amyloidoses are a large group of diseases [3] caused by proteins with altered metabolism, which have the propensity to precipitate and deposit in tissues, causing organ damage. In AL amyloidosis a plasma cell clone, usually in the bone marrow, makes the precursor, a monoclonal (MC) immunoglobulin light chain, which is released to the plasma and detectable as a free light chain component [4]. In systemic AL amyloidosis the amyloid fibrils could deposit in several organs (heart, kidney, liver, soft tissues and peripheral nervous system). The therapy of AL amyloidosis relies on a chemotherapy directed against the plasmacellular clone producing the amyloidogenic light chains [5]. Due to the multi-organ damage caused by the disease, supportive therapy (such as diuretics) for the organs involved plays a crucial role [5].

2 Methods

In this paper we focus on the representation and reasoning of effects of actions in the context of clinical guidelines. Doing so supports us in three important uses: (1) summarization and explanation of observed clinical states; (2) enhancing patient safety using safety rules; and (3) assessing guideline compliance.

We formulated our ontology using Web Ontology Language (OWL) and Protégé [6]. OWL allows us to provide definitions for the concepts relevant to the domain. To achieve consensus, we are using SNOMED-CT and LOINC codes for term IDs and relying on SNOMED's term hierarchies for lower-level abstractions, such as types of chemotherapeutic drugs.

In order to connect the ontology to patient data, our ontology has concepts that correspond to data elements from the EMR. Thus, the ontology makes explicit the data elements that are important in the domain. Moreover, our modeling is oriented toward computability as it provides computable definitions of concepts, such as risk classes, and it provides extensions of existing drug ontologies by specifying effects of clinical actions such as substance administrations.

3 The State, Action, and Effect Ontology

The main classes in the ontology include ClinicalState, Action, and EffectAction. Additional causal relationships support reasoning with the ontology. The classes and relationships are described in this section. The reasoning is described in Section 4.

3.1 Clinical States

Clinical States are abstractions over patient data. Following the Generic Process Model [7], states can be specialized into initial states (clinical characteristics at diagnosis/prior to start of therapy), goal states and intermediate process states. For example, in the domain of AL amyloidosis [5] the initial state could be that of Low risk (of death). According to the SIA therapeutic guideline [2], the LowRiskState is defined in OWL as following (the acronyms UCT and OCT used for quantitative values means "Under/Over Critical Threshold"):

LowRiskState: ClinicalState and DiagnosedWithAL-Amyloidosis and *UCT*-BNP and *UCT*-NT-proBNP and *UCT*-cTn and *OCT*-eGFR and *UCT*-ALT and *UCT*-ALP and AgeLE60

where BNP si thebrain natriuretic peptide type B; NT-proBNP-N is the terminal fragment of proBNP; cTn is cardiac troponin (type I or T); eGFR is the estimated glomerular filtration rate (by MDRD formula); ALTis Alanine Aminotransferase; ALP is Alkaline phosphatase and AgeLE60 means age less than or equal to 60.

In turn, *UCT*-BNP, a marker of cardiac dysfunction, is defined as a clinical state of a finding of an observation whose concept code corresponds to the BNP concept code from SNOMEDCT and whose quantity is ≤ 50 ng/L in healthy individuals:

UCT-BNP: ClinicalState and hasFinding some (Observation and (hasQuantity some float [≤50]) and (hasUnit value ng/L) and (concept_code value 116886006) and (terminologyName value "SNOMEDCT"))

Similarly, intermediate and high risk states can be defined.

To assess the effect of chemotherapy on the hematologic disease (direct effect on the plasma cellular clone) and on the organ damage caused by the disease (e.g., possibility of functional improvement after a hematologic response is achieved) we applied the response criteria defined by the International Society of Amyloidosis [8]. In particular, hematologic response can be "partial" or "complete". Therefore, an example of a goal state is "complete response", defined as the disappearance of the monoclonal light chains component (SNOMED code 414766003) measured by serum and urine immunofixation, and normal circulating free light chains κ/λ ratio (LOINC code 48378-4), measured by Freelite assay. This state is represented as:

CompleteResponse: ClinicalState and hasFinding some (Observation and (hasQuantity value absent) and (concept_code value 414766003) and (terminologyName value "SNOMEDCT")) and hasFinding some (Observation and (hasQuantity value normal) and (concept_code value 48378-4) and (terminologyName value "LOINC"))

An alternative way to define a ClinicalState is to define a phenomenon (i.e., an observable entity, such as serum free light chains concentration) and a quality direction (e.g., increase, decrease) that affects that phenomenon. For example:

DecreasedSerumFLC (Clinical State):
(has_quality some decrease and phenomenon some SerumFreeLightChainsConcentration)

Where SerumFreeLightChainsConcentration measurement is recorded as:

SerumFreeLightChainsConcentration:
Observation and (concept_code value 57778-3) and (terminologyName value "LOINC")

Note that according to the guideline, a state of decreased serum FLC (in contrast to a state of no MC, corresponding to complete response) is defined relative to the initial value of serum FLC. The goal state for partial response is to reduce the value to 50% of the initial serum FLC. OWL does not allow stating such temporal abstraction but it can potentially be specified in Semantic Web Rule Language (SWRL).

3.2 Actions

In order to start at the initial (or current) state and reach the goal state, the guideline recommends actions. Our ontology distinguishes among four high-level types of actions: *procedures*, such as operations or therapeutic actions (e.g., a drug administration), which affect the patient's states, *diagnostic actions*, such as laboratory tests, which allow the physician to assess the state of the patient, *complex actions*, such as a protocol which is a network of lower-level actions, and *visits*.

An example of a procedure is CyclophosphamideAdministration: CyclophosphamideAdministration is-a SubstanceAdministration whose active_principle is cyclophosphamide. Cyclophosphamide is-a AlkylatingAgent that is-a ChemotherapeuticDrug_that_AffectsCellDivisionOrDNASynthesis.

An example of a diagnostic action is VitalSignDetermination, which has-finding such as fever, blood pressure, and heart rate.

3.3 Effects of Actions

Actions of type procedures have effects. The Effect-Action class represents the *relationship* between potentially multiple actions and their effect. The effect is a clinical state caused by the actions. The properties of the Effect-Action class are:

```
Effect-Action:
  has_context ClinicalState
  has_action multiple Action
  affects multiple anatomicalEntity
  has_effect multiple ClinicalState
  desirable Boolean                    // representing the desirability of the effect
  prevented_by ActivePrinciple
```

For example, the intended effect of the combination of Cyclophosphamide, Thalidomide, and Dexamethasone (CTD) action is to achieve the goal of complete response. It is modeled in the following way:

CTDAchieveCompleteResponse:
has_context some Diagnosed_with_AL-Amyloidosis
has_action some CTDAdministration
has_effect some CompleteResponse
desireable value true

Where the clinical state CompleteResponse is defined as above.

The individual drugs also have side effects. An important side effect of Cyclophosphamide is that it could cause the patient to be in a dangerous "immunocompromised state" by decreasing neutrophils. This is modeled as follows:

CyclophosphamideResultsInImmunocompromisedState
has_action some CyclophosphamideAdministration
has_effect some NeutrophilsUnderCriticalThreshold
desirable value false

3.4 Other Causal Relationships

The examples presented above demonstrate the OWL modeling of Effect-Action relationship between an action and the ClinicalState that is a consequence of the action. Other types of relationships in our ontology include:

ClinicalState *can_cause* ClinicalState

(e.g., Immunocompromised can_cause IncreasedRiskOfInfection;
NeutrophilsUnderCriticalThreshold can_cause Immunocompromised)

While actions of type procedure have effects on clinical states and on phenomena, diagnostic actions result in findings. This is represented via the relationship:

DiagnosticAction *has_finding* Observation

(e.g., VitalSignDetermination has_finding Temperature).

Note that we can relate findings (e.g., temperature), which are outputs of diagnostic actions, to clinical states that could be inferred from certain values of these findings (e.g., a state of fever could be inferred from the result (temperature) of VitalSignDetermination). This is represented in the following relationship:

ClinicalState *phenomenon* Observation (e.g., Fever *(can be inferred from) phenomenon* Temperature)

Fig. 1 presents a semantic network depicting relationships between individuals which are members of ClinicalState, Procedure, DiagnosticAction, and Observation.

4 Reasoning with the Ontology

Following discussion of mapping the ontology to EMR data we discuss how reasoning with the ontology can support our three use-cases, while focusing on summarization and explanation of observed clinical states.

4.1 Mapping and Integration with the Electronic Patient Record

We are implementing the mapping between the ontology and the data stored in the EMR developed as part of the AMICA project [9]. This EMR is a relational database, accessible through a web interface by the physicians and the professionals involved in this project, and contains all the clinical information, measurements, observations and actions performed to treat the patients suffering from amyloidosis. All the data are stored with their timestamps, therefore it is possible to apply the reasoning process over the patient's data in a given temporal window. We highlight that this reasoning process is applied considering a specific patient at a time. Part of the retrieved information is used to calculate abstractions over the patient's data. Some of the abstractions (for example, the state abstractions obtained by checking thresholds over numerical data) can be calculated directly by the ontology applying SWRL rules. More complex abstractions (for example, trends over time) will be pre-calculated by ancillary tools of the EPR. The ontology is populated by loading the proper information with SQL queries run on the EPR. These queries are obtained from the mapping between the ontological concepts and the tables of the EPR, following the methodology described in [10, 11]. A possible mapping between the concept "temperature" shown in Figure 1 and the EPR is the following:

TEMPERATURE (X, First_date, Last_date)→SELECT TEMP, DATE FROM OBJECTIVE_EXAM WHERE PATIENT_ID=X AND DATE BETWEEN First_date AND Last_date

The raw data and abstractions can be put on a timeline, representing the succession of the patient's observations and actions during the treatment process in the considered time window. This information is used to support the reasoning processes according to the relationships defined in the ontology, as described in sections 4.2-4.4.

4.2 Summarization and Explanation of Observed Clinical States

The ontology, once integrated into the EMR, could be used to create summaries of the patient's state and the actions performed during his treatment. The ontology uses abstractions such as goal states, procedures, and diagnostic actions, which are further refined into more specific class types. These taxonomies could be used to present the patient's data according to clinical states and actions. For example:

Initial State: Low Risk Amyloidosis patient
Goal: CompleteResponse
Procedures: Goal procedure: chemotherapy (CTD)
Efficacy: AbsenceOfuMC and AbsenceOfsMC and …
Undesired effects: UCT-neutrophils on CBC
Preventive procedure for infection: antibiotics (ciprofloxacin)

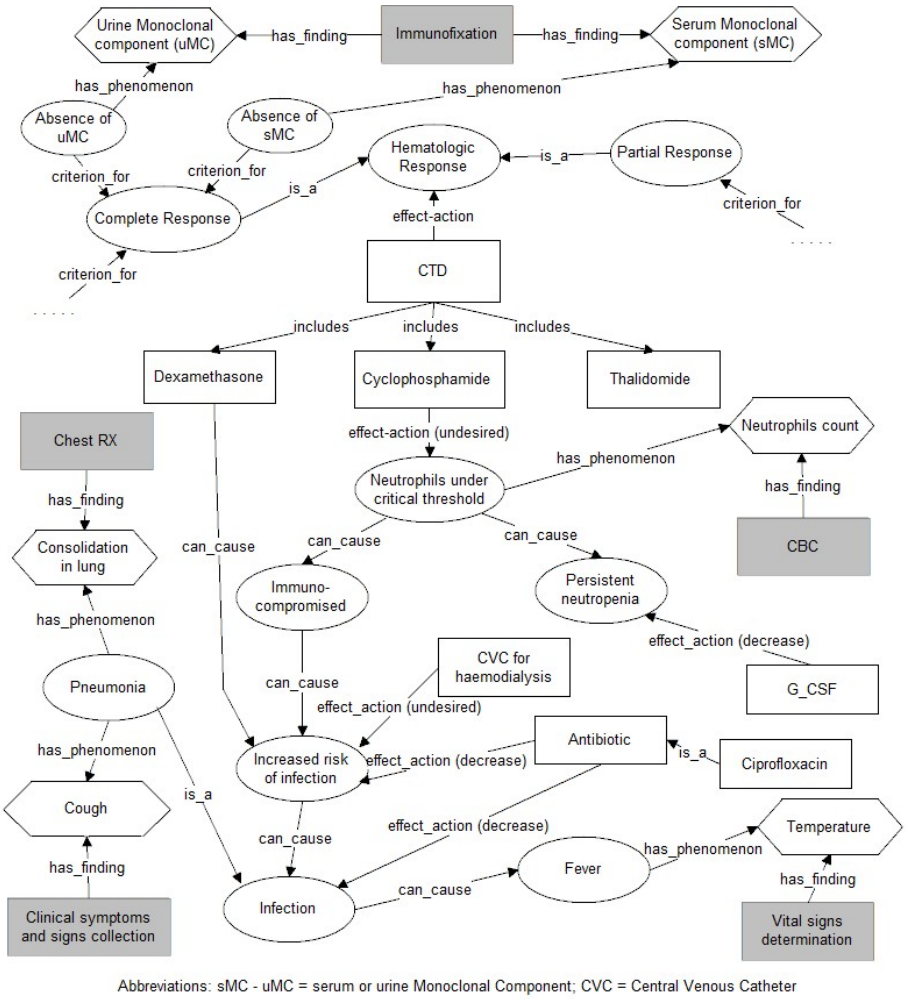

Fig. 1. A semantic network depicting relationships between ClinicalState (oval), Procedure (light rectangle), DiagnosticAction (gray rectangle), and Observation (hexagons)

The causal relationships in the ontology could connect response with treatment in a given patient context. They could be used to generate explanations of why certain actions were selected. Our intention is not just to provide explanations to physicians who are using the guideline, but to allow guideline authors to acknowledge, during design time, whether their guideline text is complete and to allow final users to find justifications for every recommendation. In fact, often the guideline authors underestimate the amount of knowledge that is not explicitly written in the guideline and that requires further reasoning to the reader to fully understand a recommendation rationale. We want to find the actions that could cause a chosen clinical state (e.g., infection), and to look for the evidence related to this state. For example, if the goal of the guideline is the CompleteResponse, then one of the possible actions who

have this desired effect is chemotherapy (CTD). This can be deduced by finding the actions which have an Effect-Action relationship with the goal state (see Fig. 1).

The recommended action may also have undesired effects. For example, as shown in Fig. 1, a chemotherapy such as cyclophosphamide causes the patient to be in a state of low neutrophils, which can cause the patient to be in an immunocompromised state. In turn, an immunocompromised state can cause a state of increased risk of infection. Thus a state of infection is an undesired state that we should be watching for when cyclophsophamide is administered, as infection could be explained by the causal chain starting with cyclophosphamide and ending in state of infection. Following the causal network of Fig. 1, we can determine how the state of infection could be monitored; a diagnosis of infection is supported (among others, not reported for lack of space) by the findings of symptoms and signs like cough and fever. These and possibly other findings supporting infection are not mentioned in the guideline text. They are part of the "common knowledge" a physician must have. Therefore, to support young doctors in reasoning, the ontology may refer to such additional knowledge.

The signs of infection, in turn, are findings that are determined by diagnostic actions. For example, fever is a finding of vital signs determination. Similarly, a state of low neutrophils is supported by a finding that is determined by the CBC (complete blood count) diagnostic action.

However, a state of infection may have other explanations (see Fig. 1). For example, a patient with AL-Amyloidosis-related renal insufficiency may need a central venous catheter (CVC) for haemodialysis, and CVC itself represents a high risk factor for infection. Therefore, the summary of the patient's state should consider whether the patient is using a catheter to suggest other causes of infection.

Since the AL-Amyloidosis guideline considers that giving chemotherapy puts the patient at increased risk of infection, the guideline suggests giving antibiotics (ciprofloxacin) with the chemotherapy cyclophosphamide if the neutrophils concentration falls below the threshold of 1500/μL, which is indicative of a state of immunocompromised patient.

An example of a patient's summary that could be generated based on the patient's data and the knowledge found in the ontology is obtained by considering the patient's information loaded from the EPR and shown in Fig. 2. The lower part of this image shows measurements (FLC, Neutrophils count, temperature) and actions (e.g. start and stop of procedures and therapies), while the upper part shows abstractions obtained using the measurements' values (normal, high, low, fever).

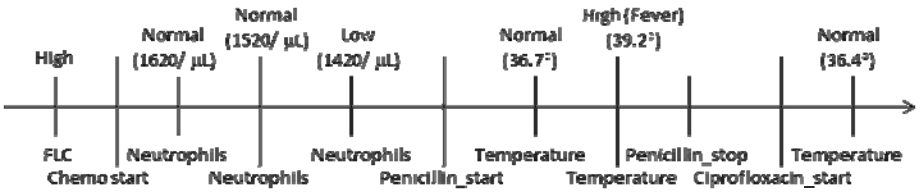

Fig. 2. Patient's information on the temporal timeline

After applying the reasoning process on the retrieved data, the following summarization can be created: "The patient was given CTD in order to lower the amyloidogenic serum free light chain concentration (represented as FLC in figure 2). His neutrophils' concentration was monitored to detect imnocompromised state as a result of the chemotherapy treatment. When the neutrophils concentration dropped below 1500 per μL, penicillin was given, apparently to prevent infection, although the recommended antibiotic is ciprofloxacin. Vital signs were followed to monitor for infection, and when fever developed, it was thought to be a consequence of the chemotherapy (the patient was not using a CVC, for example). The antibiotic was switched to ciprofloxacin and the infection signs disappeared within 24 hours".

In the following steps, we present an algorithm for reasoning with the causal relationships of the ontology, corresponding to the informal explanation exemplified above.

1. Clinical states that should be monitored are those that are the goal state or are undesired states (e.g., IncreasedSerumFLC, Infection). Undesired states are clinical states that can be undesired effect-action of actions or are reachable from such states using "can_cause" relationships. Some undesired states do not have to be monitored because they are intermediate non-measureable states that provide an explanation to other states (e.g., risk of infection does not need to be monitored as it explains the state of infection, which is monitored).
2. To monitor a clinical state, look in the patient's EMR for the evidence related to this state using *has-finding* and *can_cause* relationships (e.g., Vital_Sign_Determination *has-finding* Fever; Infection *can_cause* Fever).
3. Find the actions whose direct effect is the clinical state sought or which have a path of clinical states connecting to the clinical state sought. Connectors include the relationship "can_cause" (e.g., Chemotherapy *effect-action* UCT-neutrophils *can_cause* Immunocompromised *can_cause* Increased_Risk_of_Infection *can_cause* Infection; CVC *effect-action* Increased-Risk_of_Infection *can_cause* Infection).
4. Look for evidence for the actions found (e.g., chemotherapy was given; no CVC).

4.3 Safety Rules

At guideline authoring time or at run time, safety rule violations could be checked and generate alerts. Some examples of safety rules include:

Rule 1: If an action is performed, and this action has a potential undesired effect (either directly via the effect-action relationships or indirectly via can_cause relationships), you need to monitor for it. If the EMR doesn't record the effect directly, then follow the "can_cause" relationships and search the EMR for data for lower-level abstractions.

For example, chemotherapy could cause infection. Maybe the EMR does not record infection but records signs of infection, such as cough and fever, which are caused by infection. From these measureable signs, infection may be inferred.

Rule 2: Prevention. An undesired effect (e.g., bacterial infection) could be prevented or counter-acted by an action with a reverse effect (e.g., the death of the bacteria).

The guideline defines knowledge about prevention. For example, ciprofloxacin is given as an antibiotic to prevent a bacterial infection. Our ontology could be used to represent the fact that ciprofloxacin has the ability to kill a wide spectrum of bacteria and reasoning could be used to infer that it should be prophylactically administered when neutrophils fall below their critical threshold.

Rule 3: Two drugs that have each the same effect can reinforce each other. Thus, the effect should be monitored to see that it is not too large.

4.4 Assessing Guideline Compliance

By comparing guideline-specified actions with actual actions we can assess guideline compliance. For example, we can check whether guideline-recommended preventive actions were done (e.g., in presence of RiskOfInfection, antibiotic is a preventive action, and ciprofloxacin is an antibiotic. We can see whether the patient was given this or another antibiotic, or neither).

5 Related Work

Some of the early decision-support systems [12] reasoned with physiological explanations and mechanisms of actions. However, supporting explanations was a topic that has not been trendy in the past twenty years, due to its complexity and the fact that physicians were not resorting to explanations very often. In light of this consideration, our explanations are mainly targeted for the guideline authors, so that they could check whether their set of recommendations is completely justified. Another difference is that we used logic and not rule structures. We now review some of these early works.

In his paper from 1986, Clancey [12] discusses generation of explanations in knowledge-based systems. He classified the knowledge included in the famous MYCIN rule-based system into four types of knowledge roles: (1) The heuristic rule: A relation between data and diagnoses or therapies, (2) Structure: Subsumption relations among data, diagnoses, and therapies, (3) Strategy: The procedure for applying rules, and (4) Support: The justification or evidence for rules. In addition, Clancey added meta-rules that capture strategy, thereby separating the inference procedure from the medical knowledge. In his paper he gives the example of different types of justification for not giving tetracycline to a four-year old. If the reason is that the child is less than seven then a heuristic rule is applied. If the reason is that age is one of the contraindications then the organization of knowledge is considered. If the explanation is based on rules for when contraindications are considered and how each type is considered then strategy is used. If the explanation is that tetracycling causes chelation which results in the molecule binding to the growing teeth and bones and a social consideration that attests people don't want to have discolored teeth then support (domain) knowledge is used. While we try to capture support knowledge regarding

mechanisms of action, Clancey states: "I tried consistently to apply this analysis when working with physicians particularly to focus their explanations on strategy and avoid the bottomless pit of support explanations".

In the Digitalis Advisor, Swartout [13] researched the knowledge needed to design a DSS and how automatic programming techniques could be used to generate an expert system. He distinguished and represented domain knowledge and general strategies for solving problems. Using automatic programming techniques, he could create records of how such knowledge was applied to create the DSS. The design record captured the design of explanation: using the domain and problem solving knowledge to explain what the system was trying to say and how it was trying to say it.

In his ABEL system [14], Patil focused on the task of providing expert consultation for electrolyte and acid-base disturbances. Patil developed a multi-level representation of causal knowledge, and explored issues of the aggregation of available case specific knowledge into concise summaries of the patient's illness. Medical knowledge can be represented and reasoned using different types of knowledge: (1) anatomical knowledge, including a part-of hierarchy for organ systems, contained-in and position relations for major anatomical features, and a connected-to relation which provides material flow information, (2) physiology knowledge describes the fluid compartments of the body, the spaces of distribution of various solutes, and the relative distribution of losses and gains in the various compartments under different conditions, and (3) pathophysiology knowledge about disease etiologies, a taxonomy of disease processes, and causal relations which describe how the changes in a given state influence other states.

Pople's work on Caduceus [15] included similar causal networks used to support diagnostic reasoning. The knowledge representation included a causal network of disease states and a taxonomy of diseases. The causal network allowed starting from a finding (such as jaundice) and through intermediate states (e.g., conjugated hyperbilirubinemia) arriving at a diagnosis (e.g., duct stone, billary cirrhosis). Through this network, disease mechanism could be explained, aggregated states could be confirmed by fewer findings thereby reducing the number of hypotheses considered, and findings could be attributed to temporal disease-states. The disease taxonomy enabled associating findings with organ-system involvement. Furthermore, the taxonomy could be used to aggregate hypotheses through planning links that identify high-level nodes that subsumes all nodes causally linked to a finding. As in causal network, fewer decisions and fewer findings needed to be confirmed or refuted. This enabled quick characterization of a clinical problem by allowing combination of task definitions when several findings are involved.

Reasoning with effects of actions is similar to the notion of actions' preconditions and post-conditions that are used in the planning literature to plan a route from the current to the goal state [16]. Planning can also be used in the context of clinical guidelines; the notion of preconditions and post-conditions is part of the PROforma [17] guideline formalism. In the Asbru [18] guideline formalism, reasoning with intentions of plans is used for skeletal plan refinement during run-time, where more granular actions could be selected to refine a more abstract plan. Reasoning with pre and post conditions of actions has also been used to suggest exception-handling

mechanisms that are invoked in a goal-based manner, where exceptions raise new goals and a library is searched for medical actions that could potentially achieve those goals [19].

Our methodology for assessing guideline compliance uses abductive reasoning with causal networks to find whether physicians followed the guideline's recommendations, or used other actions with similar effects, that were appropriate for the patient's context (clinical state). The work by Advani et al. [20] on intention-based critiquing of guideline-oriented medical care also works both at the level of following guideline recommendations per se or following the intentions of the guideline (i.e., the intended effects of actions) for the patient's context. But whereas we focus on abductive reasoning with causal relationships, their work centered on finding temporal abstractions while assuming a context and a relative timepoint as a reference point. Groot et al. [21] took another temporal approach to critiquing by representing clinical guideline propositions in temporal logic, and using model checking to check whether EMR data shows adherence to these properties.

Supporting guideline authors in writing justified and safe clinical guidelines and protocols was also the focus of the Protocol Inspection and Critiquing Tool of Study Structure (PICASSO) tool [22]. This tool uses a knowledge base of medical knowledge, clinical trial protocol knowledge, and formal constraints that describe the sorts of errors we wish to find in a protocol. It evaluates a new protocol design and alerts the author if inconsistencies/errors are detected. Medical knowledge includes diseases, drug treatment, drug toxicities, drug interactions, and monitoring events. The protocol knowledgebase includes treatments, eligibility criteria, and recommended monitoring activities, Similar to our work, PICASSO links drug treatments in a protocol to potential drug toxicities, and to specific monitoring tests for detecting those toxicities. The constraints (safety rules) that PICASSO evaluates are represented in Protégé's [23] Axiom Language and include (1) do not prescribe contraindicated drugs, (2) include all required monitoring actions, and (3) do not prescribe interacting drugs.

Hammond et al. [24] derived a set of generic safety principles from examples of reasoning about the efficacy and safety of chemotherapy. Some of these principles were used in the OaSiS tool for the run time management of patients under existing chemotherapy plans. These principles were used to inhibit additions to treatment that might exacerbate a known hazard or undermine efficacy and as a filter so that drug doses. The safety principles included exacerbation, diminution, reaction, warning, monitoring, efficacy, sequencing, critiquing, and prevention. The authors demonstrated how these safety principles, represented as Prolog rules, contributed to the build time design of a chemotherapy plan.

Other researchers addressed the problem of the integration of ontologies and databases. Fankam et al. [10] proposed an extension of the ANSI/SPARC architecture for database applications, adding an ontological level for representing explicitly the semantics of data within the database. Pérez-Urbina et al. [25] describe how to answer queries over an OWL 2 QL ontology and a database via query rewriting.

6 Discussion

Our logic-based approach for modeling clinical actions and their effects allows for reasoning in many directions. First, summarization of the patient's data allows physicians to obtain an overview of the patient's current state in terms of clinical abstractions used by the guideline model. Second, safety rules can be verified on the real data and could raise alerts if safety violations occur. Third, compliance to guideline recommendations could be assessed. And fourth, our clinical guideline ontology with its causal network could be used to obtain more complete guidelines in which the recommended actions are justified in terms of mechanisms of action that yield the guideline's goals.

The algorithm that we devised for finding potential explanations for phenomena based on a causal network is a form of adductive reasoning. This reasoning process could be implemented using SWRL rules that work with the OWL-based ontology.

Our approach also has some limitations. First, in terms of expressiveness, not all of the requirements for executing guidelines can be met by an OWL specification; logic is not the most suitable approach for sequencing actions and for performing arithmetic computations, as is required for drug dosing. Therefore, in the final system, some reasoning will be done using the ontology (OWL and SWRL) and some reasoning will be done by procedural code directly embedded in the EMR interface. Another limitation is in terms of usability: we have not yet developed tools to allow users who are not experienced in logic specification to define and execute safety rules,

We have demonstrated how we can start from a clinical state that should monitored and devise a set of diagnostic actions that should be used to monitor the patient's state and preventive actions that could be used to prevent the undesired states when signs for increased risk for those states are evident. How do we know which states should be monitored? Some clinical guidelines explicitly specify the adverse effects of therapeutic actions and procedures. On other cases, this knowledge is implicit and needs to be entered into the ontology as part of the knowledge-based formulation. Making such knowledge explicit exposes missing information in the guidelines and ensures that the recommendations of the guideline are complete and justified. Once this knowledge is entered into the ontology, we can use the first step of our proposed algorithm to identify potential states that should be monitored. Our algorithm assumed that the undesired states that should be monitored are a consequence of guideline actions. However, undesired states could also be a consequence of other phenomena such as disease processes or clinical states. For example, AIDS can-cause immunocompromised. Such clinical states should be considered when providing possible explanations for patient states. Another direction that should be explored is how to rank multiple explanations according to their strength of evidence and likelihood.

6.1 Future Work

While this paper focused on the qualitative reasoning with effects of actions and its potential for three use cases, our motivation for developing this ontology was to allow rapid update of new knowledge. The new medical knowledge could be discovered by

researchers working without the ontology, or, the ontology could facilitate knowledge discovery. For example, the ontological abstractions (e.g., goal, risk group, undesired effect of action, clinical state) could be used to classify patient data in order to learn the treatments that achieved the goal state or in general, improved the patients' state, pointing to variations of therapies that were more effective or reduced harmful effects.

We would also like to explore the use of external knowledge sources such as drug knowledge bases to supply us the knowledge regarding drug effects. These knowledge sources could potentially be imported into OWL or queried without the need to replicate the external knowledge. Since the modeling of drug effects in other sources may not be the same as our modeling, we would need to reconcile the difference in a way that would be useful for our knowledge management purposes.

Lastly, to complete the summarization capabilities, we would like to explore the types of temporal abstractions that are needed for summarization and can be supported by SWRL rules.

Acknowledgments. To Mario Stefanelli, who initiated the amyloidosis project and inspired us to work on new challenges and important causes. We also thank Silvia Panzarasa and Riccardo Ferrari for fruitful discussion.

References

[1] Peleg, M., Kantor, R.: Approaches for guideline versioning using GLIF. In: Proc. AMIA Symp. 2003, pp. 509–513 (2003)

[2] Societa Italiana per l'Amiloidosi. Terapia dell'Amiloidosi AL sistemica- Linee guida (2009), `http://www.amiloidosi.it`

[3] Sipe, J.D., Benson, M.D., Buxbaum, J.N., Ikeda, S., Merlini, G., Saraiva, M.J., et al.: Amyloid fibril protein nomenclature: 2010 recommendations from the nomenclature committee of the International Society of Amyloidosis. Amyloid 17(3-4), 101–104 (2010)

[4] Westermark, P., Benson, M.D., Buxbaum, J.N., Cohen, A.S., Frangione, B., Ikeda, S., et al.: A primer of amyloid nomenclature. Amyloid 14(3), 179–183 (2007)

[5] Palladini, G., Merlini, G.: Current treatment of AL amyloidosis. Haematologica 94(8), 1044–1048 (2009)

[6] Knublauch, H., Horridge, M., Rector, A.L., Stevens, R., Drummond, N., Lord, P., et al.: The Protege OWL Experience: workshop on OWL: Experiences and Directions. In: Fourth International Semantic Web Conference, Galway, Ireland (2005)

[7] Soffer, P., Wand, Y.: Goal-driven Analysis of Process Model Validity. In: Persson, A., Stirna, J. (eds.) CAiSE 2004. LNCS, vol. 3084, pp. 521–535. Springer, Heidelberg (2004)

[8] Gertz, M.A., Comenzo, R., Falk, R.H, Fermand, J.P., Hazenberg, B.P., Hawkins, P.N., et al.: Definition of organ involvement and treatment response in immunoglobulin light chain amyloidosis (AL): a consensus opinion from the 10th International Symposium on Amyloid and Amyloidosis, Tours, France, April 18-22 (2004); Am. J. Hematol. 79(4), 319–328 (2005)

[9] Ferrari, R., Caffi, E., Rubrichi, S., Quaglini, S., Stefanelli, M., Russo, P., Palladini, G., Lavatelli, F., Merlini, G.: AMIloidosi CArtella (AMICA): an electronic patient record specifically designed for an amyloidosis network. Amyloid 17(1), 88 (2010)

[10] Fankam, C., Jean, S., Pierra, G., Bellatreche, L., Ait-ameur, Y.: Towards Connecting Database Applications to Ontologies. In: Intl. Conf. on Advances in Databases, Knowledge, and Data Applications, pp. 131–137 (2009)
[11] Horrocks, I., Sattler, U., Tobies, S.: Practical Reasoning for Expressive Description Logics. Logic J. of IGPL 8(3), 239–263 (2000)
[12] Clancey, W.J.: From Guidon to Neomycin and Heracles in Twenty Short Lessons: ORN Final Report 1979-1985. AI Magazine 7(3), 40–60 (1986)
[13] Chandrasekaran, B., Swartout, W.: Explanations in Knowledge Systems: the Role of Explicit Representation of Design Knowledge. IEEE Intelligent Systems 6(3), 47–49 (1991)
[14] Patil, R.S.: Causal Representation of Patient Illness for Electrolyte and Acid-Base Diagnosis. MIT, Cambridge (1981)
[15] Pople, H.E.: Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnostics. In: Szolovits, P. (ed.) Artificial Intelligence in Medicine. Westview Press, Boulder (1982)
[16] Fikes, R.E., Nilsson, N.J.: STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving, Technical Note 43r. Stanford Research Institute, Menlo Park (1971)
[17] Fox, J., Rahmanzadeh, A.: Disseminating medical knowledge: the PROforma approach. Artificial Intelligence in Medicine 14, 157–181 (1998)
[18] Shahar, Y., Miksch, S., Johnson, P.: An Intention-Based Language for Representing Clinical Guidelines. In: Cimino, J.J. (ed.) AMIA Annual Fall Symposium, pp. 592–596. Hanley & Belfus, Washington, D.C (1996)
[19] Grando, A.M., Peleg, M., Glasspool, D.: A goal-oriented framework for specifying clinical guidelines and handling medical errors. J. of Biomedical Informatics 1(2), 287–299 (2010)
[20] Advani, A., Lo, K., Shahar, Y.: Intention-Based Critiquing of Guideline-Oriented Medical Care. In: Proc. AMIA Annual Symposium, pp. 483–487 (1998)
[21] Groot, P., Hommersom, A., Lucas, P.J., Merk, R.J., ten Teije, A., van Harmelen, F., et al.: Using model checking for critiquing based on clinical guidelines. Artif. Intell. Med. 46(1), 19–36 (2009)
[22] Rubin, D.I., Gennari, J., Musen, M.A.: Knowledge representation and tool support for critiquing clinical trial protocols. In: Proc. AMIA Symp., pp. 724–728 (2000)
[23] Gennari, J., Musen, M.A., Fergerson, R.W., Grosso, W.E., Crubezy, M., Eriksson, H., et al.: The Evolution of Protege: An Environment for Knowledge-Based Systems Development. International Journal of Human-Computer Interaction 58(1), 89–123 (2002)
[24] Hammond, P., Modgil, S., Wyatt, J.C.: Safety and computer-aided design of chemotherapy plans. Top Health Inf. Manage 20(4), 55–66 (2000)
[25] Pérez-Urbina, H., Horrocks, I., Motik, B.: Efficient Query Answering for OWL 2. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) ISWC 2009. LNCS, vol. 5823, pp. 489–504. Springer, Heidelberg (2009)