

[Click here to view linked References](#)

Noname manuscript No. (will be inserted by the editor)
--

Decision Support Provided by a Temporally Oriented Health Care Assistant

An Implementation of Computer-Interpretable Guidelines

Tiago Oliveira · António Silva · José
Neves · Paulo Novais

Received: date / Accepted: date

Abstract The automatic interpretation of clinical recommendations is a difficult task, even more so when it involves the processing of complex temporal constraints. In order to address this issue, a web-based system is presented herein. Its underlying model provides a comprehensive representation of temporal constraints in Clinical Practice Guidelines. The expressiveness and range of the model are shown through a case study featuring a Clinical Practice Guideline for the diagnosis and management of colon cancer. The proposed model was sufficient to represent the temporal constraints in the guideline, especially those that defined periodic events and placed temporal constraints on the assessment of patient states. The web-based tool acts as a health care assistant to health care professionals, combining the roles of focusing attention and providing patient-specific advice.

Keywords Clinical Decision Support · Computer-Interpretable Guidelines · Ontologies · Temporal Representation · Reminder Systems

1 Introduction

Although exhaustive proof of the advantages of the widespread use of Clinical Decision Support Systems (CDSSs) [Musen et al, 2006] is lacking, there is isolated evidence of outcome improvements brought on by these systems in specific settings [Kaushal et al, 2003; Adhikari et al, 2005; Bright et al, 2012]. Current challenges in CDSS development are mainly concerned with making these systems user-centric and easily accessible by prioritizing and filtering the recommendations that are presented to users at a given time and place [Sittig et al, 2008].

T. Oliveira · A. Silva · J. Neves · P. Novais
University of Minho, Braga, Portugal
E-mail: {toliveira,jneves,pjon}@di.uminho.pt, pg25307@alunos.uminho.pt

1 A way to answer the challenges presented to CDSSs is to create func-
2 tionalities that enable health care professionals to track and follow up their
3 patients, schedule clinical procedures that should be performed, and manage
4 the temporal constraints placed on those procedures. Machine-readable ver-
5 sions of Clinical Practice Guidelines (CPGs) make the answer to these chal-
6 lenges possible. The use of Computer-Interpretable Guidelines (CIGs) [Isern
7 and Moreno, 2008; Peleg, 2013] endows systems with the capability of pro-
8 viding decision support across different clinical domains and situations, from
9 diagnosis to treatment, determining what questions to ask, tests to perform,
10 the value of results, and paths to follow.

11 The work presented herein tackles a gap in the development of CIGs, which
12 is the absence of efforts to include a systematic approach for embedding tem-
13 poral constraints and capabilities in CIG CDSSs Peleg [2013]. Such features
14 are important when CIGs are deployed in sensitive contexts, like those involv-
15 ing the prescription of different drug regimens that must be taken at specific
16 times, or the enactment of therapies at particular moments. These are situa-
17 tions that cannot be overlooked during medical practice as their mismanage-
18 ment may bring drastic consequences upon patients. Therefore, tools to detect
19 errors and monitor their execution actively contribute to the improvement of
20 care. Taking this into consideration, the contributions of this work are twofold:

- 21 – The first is a temporal representation model for CIG clinical tasks that
22 allows for the expression of a variety of temporal patterns, including dura-
23 tions, periodicities, task scheduling, time delays and the temporal as-
24 sessment of patient conditions. To demonstrate the expressiveness of the
25 model, a short example featuring a CPG is provided;
- 26 – The second is a form of implementation of CIG temporal execution that
27 uses the developed temporal representation. This is provided as a tool that
28 builds an agenda for the health care professional with the activities that
29 he has to perform. The tool schedules the execution of clinical tasks and
30 keeps track of their execution while trying to promote the fulfilment of
31 their temporal constraints.

32 The underlying model for CIGs is formalized in Web Ontology Language
33 (OWL) [McGuinness and Van Harmelen, 2004], with a particular focus on the
34 temporal representation of clinical tasks, which is crucial for the automatic
35 interpretation of clinical recommendations and their integration in the daily
36 practice of health care professionals. The functionalities of the system are
37 delivered through a web application posing as a health care assistant that pro-
38 vides recommendations for handling a patient, controls their execution times,
39 and provides notifications of their temporal landmarks. In terms of roles, the
40 aim is to develop a tool that focuses the attention of health care professionals
41 and provides patient-specific recommendations.

42 This article is organized in four sections. Section 2 contains a description
43 of the main existing models and tools for the temporal representation and
44 execution of CIGs, their strengths, and their limitations. Section 3 presents the
45 architecture of the system and its temporal model for CIG recommendations. It

1 also describes a case study used to demonstrate the expressiveness of the model
2 and the approach followed to make CPGs represented according to it available
3 for execution. Section 4 presents conclusions about the work developed so far
4 and future work considerations.
5
6

7 8 **2 Related Work** 9

10 The temporal constraints in CPGs are used to express a variety of elements
11 that need to be controlled in order to ensure the correct application of rec-
12 ommendations and the proper management of patients. Their correct inter-
13 pretation is vital for the integration of CPG recommendations in the practice
14 of health care professionals. In this regard, it was possible to identify two
15 main groups of temporal constraints [Shahar et al, 1998; Vollebregt et al,
16 1999; Boxwala et al, 2004; Anselma et al, 2006; Samwald et al, 2012]. The
17 first group includes temporal constraints about the execution of clinical tasks,
18 which determine when tasks should start and end. The following temporal
19 patterns are featured in this group:
20

- 21 – *Durations*: restrictions that specify for how long a task should be executed;
- 22 – *Repetitions*: restrictions that specify how many times a task should be
23 executed;
- 24 – *Periodicities*: restrictions that specify how often a task should be executed
25 and the time interval between executions;
- 26 – *Waiting Times*: restrictions that specify how long it is necessary to wait
27 between the ending of a previous task and the start of a new task;
- 28 – *Repetition Conditions*: restrictions that specify conditions regarding the
29 state of the patient that must hold true before the repetition of a task.
30
31

32 Despite being closely related, the concepts of repetition and periodicity
33 convey different meanings. In the context of CIG representation and execution,
34 a repetition only specifies how many times a task should be executed, with no
35 reference to the time between executions, whereas a periodicity specifies how
36 regularly a task execution should take place. A periodicity is usually bound
37 by a limit to the execution of a task, which may be expressed as a repetition
38 value, the same is to say a task should be performed with a specific frequency,
39 a certain number of times.
40

41 The second group encompasses temporal constraints about the state of the
42 patient. They are used to specify the temporal horizon over which a patient
43 will manifest, or should have manifested, a health state. In this sense, they
44 may be used to reason about the past or the future of the patient.

45 Table 1 shows an assessment of the most prominent models for CIG repre-
46 sentation with regards to the above mentioned groups of temporal restrictions.
47 Except for Arden Syntax [Samwald et al, 2012], which provides representation
48 primitives only for one clinical recommendation, all the other models allow the
49 definition of networks of clinical tasks.
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1 Arden Syntax Samwald et al [2012] represents clinical recommendations
2 as independent modules, called Medical Logic Modules (MLMs), comprising
3 relevant knowledge for only one decision step. In this sense, a whole CPG
4 cannot be represented in a MLM instance, but rather in a set of isolated
5 MLMs. This is regarded as one of the major limitations of this CIG model.
6 Each MLM is an ACII file divided in three partitions: maintenance, library and
7 knowledge. The actual clinical recommendation is contained in the knowledge
8 partition, where, besides knowledge slots for the definition of actions and rules
9 involving clinical parameters, there are slots for the definition of temporal
10 constraints to be applied to the clinical recommendation. However, there are
11 temporal slots only for durations and waiting times of actions. As shown in
12 Table 1, all of the other above-mentioned temporal constraints are absent from
13 this model, which indicates that Arden Syntax presents important limitations
14 in this regard.
15

16 The Guideline Interchange Format (GLIF3) [Boxwala et al, 2004] follows a
17 task network model in which a CPG contains different steps, namely: decision
18 steps, patient state steps, branch steps, synchronization steps or action steps.
19 This approach was developed in order to build flowcharts of clinical proce-
20 dures. In terms of temporal constraints, GLIF3 only allows the definition of
21 durations for actions and decisions, leaving out the definition of repetitions,
22 periodicities, waiting times, and repetition conditions. Despite this, among the
23 selected models of Table 1, it is the only one that allows for the definition of
24 temporal constraints about the state of a patient, on patient data, to evalu-
25 ate their occurrence. Yet, these temporal constraints can only be defined for
26 the verification of patient states that have occurred retrospectively. It is not
27 possible to define that a certain state should be verified in the future, within
28 a certain time.
29

30 A CIG model that offers a comprehensive representation of this pattern is
31 Asbru [Shahar et al, 1998]. This formalism is based on the notion of plan, which
32 represents a CPG. The knowledge required to perform a plan is defined by its
33 knowledge roles, which include preferences, intentions, conditions, effects and
34 plan body. The content of a plan body is composed of other plans until they are
35 no longer decomposable and stand for activities. Asbru temporal annotations
36 allow the definition of durations, with additional specification of starting and
37 ending points of a CIG activity. Furthermore, this can be done with a degree
38 of uncertainty, as these time points can be expressed as intervals. An activity
39 within a plan can also be bounded by repetitions, periodicities and waiting
40 times. As described in Table 1, it is not possible, however, to define repetition
41 conditions or temporal constraints about the state of a patient.
42

43 In the PROforma model [Vollebregt et al, 1999], CPGs are modelled as
44 plans. Each one contains atomic tasks such as actions, decisions and enquiries.
45 The temporal constraints present in this CIG model allow the specification
46 of task duration, waiting times between tasks and repetition conditions of ac-
47 tions. PROforma does not allow for the definition of periodicities and temporal
48 constraints about the state of a patient.
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

As seen in Table 1 the GuideLine Acquisition, Representation and Execution (GLARE) [Anselma et al, 2006] model is the most complete. GLARE defines a proprietary graph-based structure, where a clinical action is represented by a node. It is possible to define atomic actions like queries to obtain information, work actions that represent medical procedures, decision actions with sets of conditions, and conclusions that describe the output of a decision. Besides the representation elements for the definition of durations, repetitions, waiting times, and repetition conditions, this model is specialized in the representation of periodic actions, offering a wide array of constructs for this temporal pattern. Its only drawback is the absence of temporal constraints about the state of a patient.

It is possible to see that each model has at least one limitation in one type of temporal constraint. From that, one may conclude that the efforts in defining new CIG models have not been coupled with a simultaneous effort to devise comprehensive temporal constructs in said models.

A crucial component to the operationalization of CIGs is an execution engine that interprets the knowledge formalized in a given model and is capable of making inferences upon it, and a tool to deliver those inferences to health care professionals in the form of recommendations. The last should also present these recommendations to the users and enable inputs to feed the inference process of the execution engine. Examples of such tools include the Guideline Execution Engine (GLEE) [Wang et al, 2004], SAGEDesktop [Berg et al, 2004], or the execution engine of GLARE [Terenziani et al, 2004]. However, these tools are limited and focus mainly on displaying CPGs as oriented graphs, with no means of integration of the recommendations provided by CPGs in the daily schedule of health care professionals Isern and Moreno [2008].

Table 1 Assessment of CIG models. The ✓ indicates that the model fully represents the temporal constraint in question and the ✗ indicates the model does not represent it or has limitations in representing it.

CIG Model	Temporal constraints about the execution of tasks					Temporal constraints about the state of a patient
	Durations	Repetitions	Periodicities	Waiting Times	Repetition Conditions	
Arden Syntax [Samwald et al, 2012]	✓	✗	✗	✓	✗	✗
GLIF3 [Boxwala et al, 2004]	✓	✗	✗	✗	✗	✓
Asbru [Shahar et al, 1998]	✓	✓	✓	✓	✗	✗
PROforma [Vollebregt et al, 1999]	✓	✓	✗	✓	✓	✗
GLARE [Anselma et al, 2006]	✓	✓	✓	✓	✓	✗

3 Development of the Health Care Assistant

The solution proposed for the challenges mentioned in Section 1, i.e, tailoring the recommendations of CDSSs to the context (namely to the temporal dimension), is supported by an architecture such as the one of the CompGuide

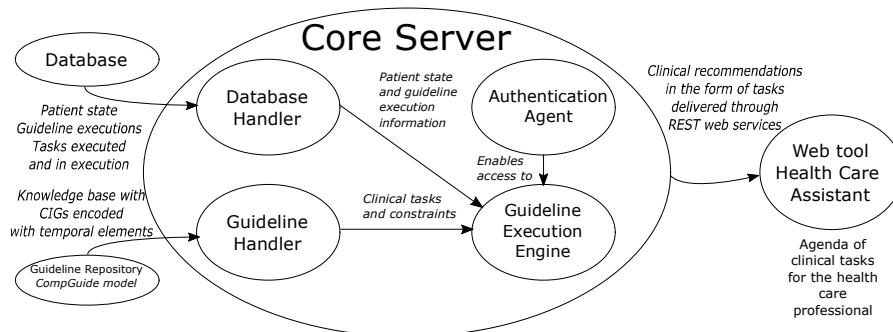
1 system, represented in Fig. 1. The architecture brings together all the ele-
 2 ments that make possible patient tracking, patient follow-up, scheduling of
 3 procedures, and monitoring of procedure constraints based on the temporal
 4 elements of CIG clinical tasks. A representation model for temporal constraints
 5 plays an important part in the whole work-flow of CIG deployment and ex-
 6 ecution.
 7

8 Fig. 1 gathers a set of elements aimed at providing timely CPG advice to
 9 health care professionals. As a system, it assumes the role of a reminder tool
 10 for focusing attention and producing patient-specific advice, by means of its
 11 end user application, the Health Care Assistant (HCA).

12 Section 3.1 describes each element of the architecture and how they con-
 13 nect to each other. Then, Section 3.2 focuses on the CIG representation model
 14 used for the CPGs in the *Guideline Repository*. After this explanation, Sec-
 15 tion 3.3 explains the temporal elements of the model, which determine how
 16 recommendations are interpreted by the execution engine and have influence
 17 in the way they are presented to a health care professional. This presentation
 18 is dealt with in Section 3.5, which provides a description of the HCA tool for
 19 the visualization and execution of CIG recommendations.
 20
 21
 22

23 3.1 Architecture of the Supporting System

24 The architecture of the CompGuide system is shown in Fig. 1. Its main com-
 25 ponent is a *Core Server* that encapsulates the most important modules of
 26 the system. The *Core Server* provides all the required services to allow exter-
 27 nal applications, such as web applications or mobile applications, to execute
 28 guidelines.
 29
 30
 31



32
33
34
35
36
37
38
39
40
41
42
43
44 **Fig. 1** Architecture of the CompGuide and modules in the *Core Server*.

45
46
47 The *Authentication Agent* is responsible for authenticating and authorizing
 48 the user to access the services of the system and, thus, allowing the access to
 49 the functionalities of the *Execution Engine*. It makes distinctions between two
 50
 51
 52
 53
 54
 55
 56
 57
 58
 59
 60
 61
 62
 63
 64
 65

1 types of users, those who can only manipulate information about guideline ex-
2 ecutions, *simple users*, and those who, in addition, can manipulate information
3 about other users, *admins*.

4 The required methods to manage and process data about patient profiles,
5 patient states, guideline executions, and tasks to be applied or currently being
6 applied are defined in the *Database Handler*.

7 The system's knowledge base, i.e., the CPGs encoded in a machine readable
8 format, are in a *Guideline Repository* accessed through a *Guideline Handler*
9 module using the OWL API. This module provides the clinical tasks and
10 respective constraints to the *Guideline Execution Engine* for interpretation.
11 The model used for CIG representation is the CompGuide model, based on
12 OWL. As described further ahead, this model incorporates temporal elements
13 for the different temporal patterns that allow for the definition of constraints
14 in clinical tasks.

15 The *Guideline Execution Engine* performs verifications on task ordering
16 and task constraints by comparing the guideline careflow with the state of the
17 patient. The result is a recommendation in the form of the next clinical task
18 to be applied. The constraints, including temporal constraints, are defined
19 directly in the ontology. Semantic Web Rule Language (SWRL) is not used
20 for this specification due to the flexibility and complexity required for this
21 definition.

22 The *Core Server* makes the functionalities of the *Guideline Execution En-*
23 *gine* available through a set of RESTful web services for: next task calculation,
24 verification of pending guideline executions, and editing of patient informa-
25 tion. The *Core Server* is implemented in Java, using the REStEasy API over
26 a WildFly Application Server. The notion of CPGs as services, present in
27 CompGuide Oliveira et al [2014], aims to facilitate the integration of CIGs
28 into any type of application and make them widely accessible, thus enabling
29 differently oriented implementations. The HCA is one such implementation.

32 33 34 3.2 Representation of Computer-Interpretable Guidelines

35 The *Guideline Repository* contains instances of different CPGs in a machine-
36 readable format. The model used is the CompGuide ontology Oliveira et al
37 [2013], which provides representation primitives for clinical recommendations
38 based on OWL. It follows a task network model in which each recommendation
39 assumes the form of a task. The tasks are connected to each other to build a
40 work-flow of clinical procedures. Although the focus of this paper is placed on
41 the temporal aspects of CIGs, it is important to present the basic structure of
42 a CPG in CompGuide. The key task classes are subclasses of *ClinicalTask*, as
43 can be seen in Fig. 2. They include the following:

- 44 – *Action*: a task that should be performed by a health care professional such
45 as an observation, procedure, exam, or treatment application;
 - 46 – *Question*: a task to get information about the clinical parameters that
47 build the state of the patient;
- 48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1 – *Decision*: a task that encodes a decision regarding the state of a patient;
- 2 – *Plan*: a composed task containing instances of the other tasks, defined to
- 3 achieve a specific goal.

4
5 The relative order of clinical tasks in CompGuide is defined with object
6 properties connecting task instances. In this regard, it is possible to define
7 sequential tasks, parallel tasks which should be executed simultaneously, and
8 alternative tasks from which one is automatically selected for execution. In
9 this sense, a guideline in CompGuide resembles a linked list of recommendations.
10 Additionally, it is possible to define different types of conditions that
11 constrain task execution, including trigger conditions to select one amongst al-
12 ternative tasks, pre-conditions which must be verified before executing a task,
13 and expected outcomes for clinical tasks. The *Condition* class allows the rep-
14 resentation of these conditions with specific properties for clinical parameters
15 and their values.

18 3.3 Temporal Elements of Guidelines

19
20 Fig. 2 shows a diagram of the classes representing temporal elements in the on-
21 tology and their relationship with non-temporal classes such as clinical tasks
22 and conditions. The classes that enable the representation of temporal re-
23 strictions are all subclasses *TemporalElement* Silva et al [2015]. One of those
24 subclasses is *TemporalUnit* which represents the different units in which a tem-
25 poral constraint may be expressed. It is an enumerated class consisting of the
26 instances *second*, *minute*, *hour*, *day*, *week*, *month*, and *year*. The remaining
27 classes enable the definition of temporal restrictions about the execution of
28 tasks and temporal constraints about the state of a patient.

31 3.3.1 Temporal Constraints about the Execution of Clinical Tasks

32
33 The *Duration* class enables the definition of how long *Actions* and *Plans* should
34 last, since these are the only tasks that may unfold continuously over time.
35 A task instance is connected to a *Duration* instance through the *hasDuration*
36 object property. There are two ways of defining *Duration* instances, as shown
37 in Fig. 3. The first is defining a minimal and maximal duration with the data
38 properties *minDurationValue* and *maxDurationValue*, which contain numerical
39 decimal values. The alternative is to define a fixed duration for the clinical
40 task with the property *exactDurationValue*. Within a *Duration* instance these
41 properties are associated with a *TemporalUnit* through the *hasTemporalUnit*
42 object property, which connects them to one of the above-mentioned instances
43 of the class. Regarding the interpretation of *Duration*, when an *Action* or a
44 *Plan* with this temporal pattern is selected for execution by the *Execution*
45 *Engine*, the HCA determines its temporal landmarks, i.e., its starting point
46 and ending point(s).

47
48 Often times there are instructions in a CPG to delay a procedure in order to
49 observe the evolution of a patient. In the CompGuide ontology this is expressed

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

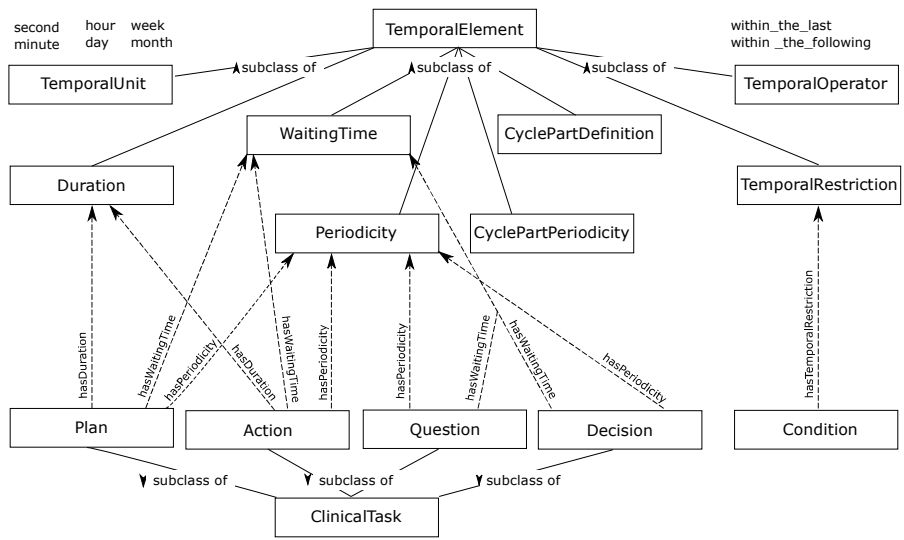


Fig. 2 Classes of the temporal model used in CompGuide.

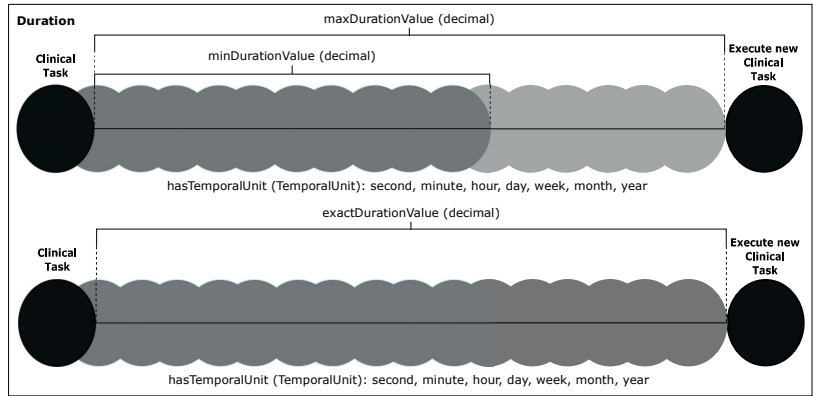


Fig. 3 Representation of a *Duration* applied to a clinical task.

with an instance of the *WaitingTime* class, by connecting the clinical task that should be delayed to the instance through the *hasWaitingTime* object property. These delays can be defined for any type of task. In Fig. 4, it is shown that the *minWaitingTimeValue* and *maxWaitingTimeValue* data properties are used when one aims to express the earliest and latest possible starting points of the task, after a previous task is finished. If the delay is a fixed value, then it is expressed with the *exactWaitingTimeValue*. The *hasTemporalUnit* property is used again to specify the units. The temporal landmarks produced by the HCA upon the interpretation of this task consist of its possible starting points.

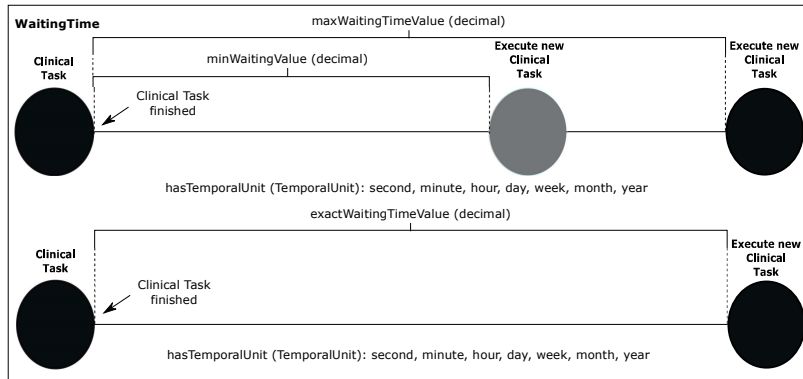


Fig. 4 Representation of a *WaitingTime* applied to a clinical task.

A periodic task is defined using the property *hasPeriodicity*, which connects the task to an instance of the class *Periodicity*. This temporal pattern can be defined for any type of task. As shown in Fig. 5, an instance of *Periodicity* can also be connected to an instance of *Duration* through the *hasDuration* object property, thus determining for how long a periodic task should take place. If one wants to state the number of times the event should take place, i.e., the number of cycles of the periodic task), it is necessary to formulate a repetition constraint, which is possible with the *repetitionValue* data property, with a range of integer numerical values. It could also be the case the periodic task should only occur until a condition about the state of a patient is verified. To express this, one uses the *hasStopCondition* object property to connect an instance of *Periodicity* to instances of the class *Condition*. While it is possible for a periodicity to have a duration and a stop condition, a repetition value and a stop condition, or just a stop condition, it is not possible to have both a duration and a repetition value because it is considered to be redundant information. The stop condition takes precedence over the other temporal restrictions, and, if the condition is met, the task is immediately stopped. The frequency of the event is defined in the data property *periodicityValue* and the associated *TemporalUnit*. A periodic task is thus unfolded in a series of executions handled as events. In turn, each event may have an associated periodicity or duration. These nested temporal patterns are defined with the *hasCyclePartDefinition* object property, connecting the *Periodicity* instance to a *CyclePartDefinition* instance, within which it is possible to define a duration with the reuse of the *Duration* class or a new periodicity with the *CyclePartPeriodicity* class. The temporal landmarks produced by the HCA for these temporal constraints consist of every execution, with starting and ending points, of the events of the periodic task.

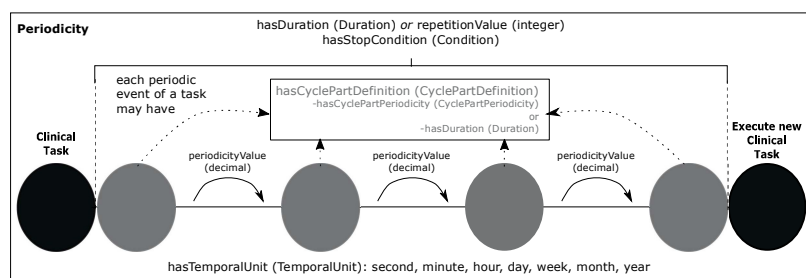


Fig. 5 Representation of a *Periodicity* applied to a clinical task.

3.3.2 Temporal Constraints about the State of a Patient

Temporal reasoning about the state of a patient is enabled by the *TemporalRestriction* class, whose instances can be associated with a *Condition* through the *hasTemporalRestriction* property. With the *hasTemporalOperator* property a *TemporalOperator* is specified for the restriction.

Being an enumerated class, *TemporalOperator* consists of two instances, *within_the_last* and *within_the_following*. The operator *within_the_last* is used when one wants to express that a condition about the patient state must have held true at least once, within a period of time just before execution time. It is used in trigger conditions, pre-conditions and conditions of rules in *Decision* instances. The execution engine interprets this operator by checking if, in the state of the patient, there is a record regarding the parameter in the condition, registered within the specified time frame, whose value validates the condition. As for the *within_the_following* operator, it expresses a condition about the future, in which one aims to observe the effect a clinical task has after being applied to a patient. Such conditions are used in task outcomes.

The temporal operators represent the reach of a temporal constraint and are used together with temporal units, defined through the *hasTemporalUnit* object property. The temporal restriction values are expressed through data properties such as *maxTemporalRestrictionValue* and *minTemporalRestrictionValue* for an interval, or *temporalRestrictionValue* for an exact value, with a range of decimal numerical values.

3.4 Case-study featuring a Guideline for Colon Cancer Treatment

In order to verify the expressiveness of the temporal representation model, a guideline of the National Comprehensive Cancer Network (NCCN) [Benson et al, 2013] was fully represented using the CompGuide ontology. The CPG is used for the diagnosis and management of colon cancer and, thus, contains many and varied clinical tasks with temporal constraints. The process of representing the guideline was accomplished using Protégé and resulted in an ontology *owl* file with 223 task instances, of which a large majority (190)

1 consisted of *Action* tasks. Among the clinical tasks, 95 of them had temporal
2 constraints. The most common type of temporal constraint was the *Periodic-*
3 *ity*, featured in 79 tasks, most of them limited by a duration. There were also 7
4 tasks with nested periodicities using *CyclePartDefinition*. The reason for such
5 an abundance of periodicities is the detailed descriptions of chemotherapy reg-
6 imens in the CPG. The remaining temporal restriction cases were 7 instances
7 of *Duration* and 2 instances of *WaitingTime*. The temporal classes and their
8 respective properties enabled the representation of all the temporal patterns
9 in the CPG. Figs. 6 and 7 show the instantiation of case examples for each
10 temporal pattern.
11
12
13
14
15
16
17
18
19

20 As can be seen in the figures, the duration of Case 1 is expressed with an
21 interval and the waiting time of Case 2 is expressed with an exact value. The
22 interpretation of the HCA in Case 1 would be to establish in the calendar
23 of the health care professional (the user) the starting and ending times of
24 neoadjuvant therapy, notifying him of when the task should start and when it
25 reaches its earliest ending time and latest ending time. As for Case 2, the HCA
26 would not let the reevaluation start right after the ending of chemotherapy and
27 would notify the health care professional of when the task should start, i.e.,
28 after 2 months.
29
30
31
32
33
34
35

36 Cases 3, 4, and 5 from Fig. 6 represent situations of *Periodicity*. In Case 3
37 the periodicity of the physical exam is bounded by a duration, which means
38 that the HCA would tell the user every 6 months during 2 years that he should
39 perform the exam. This represents the unfolding of a clinical task into multiple
40 occurrences to which we call events. The difference to Case 4 is that the last
41 is also bounded by a stop condition. Upon the ending of each colonoscopy
42 event, the HCA would ask the health care professional if signs of adenoma
43 were found and, if that were the case, it would finish the task and recommend
44 the next procedure. Case 5 has a nested periodicity that is interpreted by the
45 HCA in the following way: besides notifying the user every 3 months of the
46 chemotherapy, within each event, it would alert the user to the administration
47 of chemotherapy substances every 12 hours during 14 days. The 3 months to
48 the next chemotherapy event would start counting again after those 14 days.
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Cases	Instantiations
Case 6 "the tumor should become operable after 6 months of FOLFOX or CapeOx chemotherapy"	<ul style="list-style-type: none"> Action: chemotherapy with FOLFOX or CapeOX Condition <ul style="list-style-type: none"> parameter: tumor status value: operable TemporalRestriction <ul style="list-style-type: none"> temporalRestrictionValue: 6.0 TemporalUnit: month TemporalOperator: within_the_following
Case 7 "for therapy after third progression consider experimental chemotherapy, if the regorafenib regimen has been applied within the last 12 months"	<ul style="list-style-type: none"> Action: experimental chemotherapy after first progression Condition <ul style="list-style-type: none"> parameter: applied regorafenib chemotherapy value: yes TemporalRestriction <ul style="list-style-type: none"> temporalRestrictionValue: 12.0 TemporalUnit: month TemporalOperator: within_the_last

Fig. 7 Instantiation of case examples for temporal constraints about the state of the patient.

In the representation of the CPG there were only 6 occurrences of temporal constraints about the state of a patient. Case 6 represents the typical situation of expressing the outcome of a chemotherapy task. In this case, the HCA, following 6 months from the end of chemotherapy, would ask the user if the tumor became operable and the objective of the task was fulfilled. Depending on the answer, different procedures would be selected according to the CPG careflow. Thus, it is a condition defined for the future of the patient. As for Case 7, it configures a situation of a trigger condition for the selection of a chemotherapy regimen which has an associated temporal restriction, in order to avoid conflicts with different chemotherapy regimens. In such a case, the HCA verifies if there is a regorafenib chemotherapy regimen in the patient record within 12 months prior to execution time. Only if that were the case, would the experimental chemotherapy be selected.

In terms of expressiveness, this approach goes beyond models like Arden Syntax and GLIF3 since it provides a wider set of different temporal constraints. Taking into account what is presented in Section 2, Arden Syntax enables the definition of durations for decisions, which the presented temporal elements of CompGuide do not. However, this was a design decision as it was considered that only actions can unfold over time and, thus, have a duration. The possibility of defining temporal constraints about the state of the patient and temporally reason about patient data is the major strength of GLIF3. The CompGuide model also possesses such temporal constraints, with the additional feature of defining temporal constraints for the outcomes of actions, and not only retrospectively. Moving on to the representation of repetitions, periodicities and waiting times, the CompGuide temporal elements are on par

1 with Asbru, PROforma and GLARE. The treatment of periodicities is the
2 most important aspect in view of the frequency with which this temporal pat-
3 tern appears in CPGs. In this regard, the handling of nested periodicities is
4 similar to the one performed in GLARE. An important difference in inter-
5 pretation is that of repetition conditions. In the proposed model, rather than
6 defining a condition for the continuous execution of a task (as in GLARE),
7 the choice was to define a condition which, when true, stops the execution
8 of a task. Nonetheless, the practical effects of this are the same. The distinc-
9 tive feature of the CompGuide temporal elements compared to GLARE is the
10 definition of temporal constraints about the state of a patient.

11 The objective of the temporal elements described in Section 3.3 and exem-
12 plified in the current section is to provide an encompassing representation of
13 temporal patterns, without the need for proficiency in a specific programming
14 language for their definition. For this reason, the basic structure of these ele-
15 ments is defined in OWL, but all the meaning behind them and the logics of
16 their interpretation are encoded in the *Guideline Execution Engine*. The exe-
17 cution engine is tailored to the model and, as a result, the two should coexist
18 in a system. The implementation of the temporal model in another CDSS has
19 to be coupled with the *Guideline Execution Engine*. This is one of the reasons
20 for the creation of the *Core Server* of Fig. 1. It exposes the functionalities of
21 the *Guideline Execution Engine* through a set of web services for integration
22 in external software tools. Provided that the intended CPG is encoded accord-
23 ing to the CompGuide model in the *Guideline Repository*, its execution in an
24 external CDSS is possible.

25 3.5 Web-based Tool for the Visualization and Execution of Guidelines

26 The way in which clinical recommendations are delivered to health care pro-
27 fessionals may dictate the adoption of a tool for clinical decision support. The
28 temporal elements in the CompGuide ontology enable not only the temporal
29 execution of CPGs but also the development of new ways to visualize CPG
30 advice. Ways that allow health care professionals to accurately track their ac-
31 tivities while benefiting from automatic reasoning features, according to the
32 clinical constraints defined in a CIG. The *Guideline Execution Engine* inter-
33 prets the content of a CIG and the generated information is passed on to the
34 HCA, which is the tool through which the health care professional interacts
35 with the advice.

36 The HCA was developed as a web application so that it can be widely
37 available, whichever the platform it is accessed from. Its main objectives are
38 to provide timely clinical recommendations and integrate them in the clinical
39 practice of the health care professional. To fulfil this, it implements the func-
40 tionalities available in the *Execution Engine*. It was developed following the
41 Model-View-Control (MVC) paradigm using Java Server Faces (JSF).

42 Besides the automatic calculation of the proper clinical tasks to apply and
43 the validation of conditions regarding the state of the patient placed upon
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

them, based on user inputs, its strength lies in its temporal features. The tool builds a schedule for the health care professional based on the tasks recommended by the *Execution Engine* and their respective temporal constraints, which can be viewed as calendar, as in Fig. 8, or as a timeline, as in Fig. 9. The application informs the users of when they should execute clinical tasks, when they should start them, when they should finish them, and assesses results of expected outcomes.

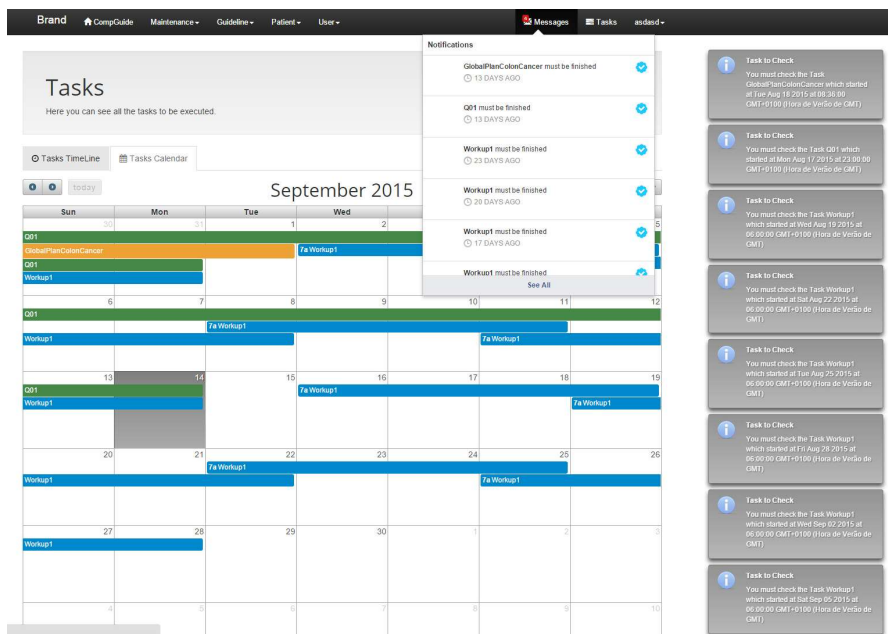


Fig. 8 Calendar task view and notifications of the Health Care Assistant.

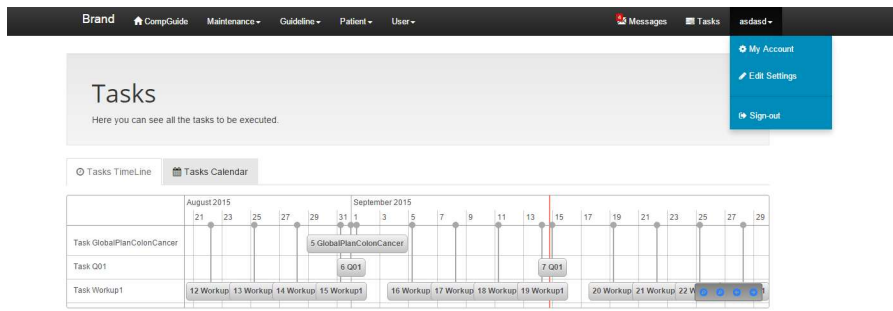


Fig. 9 Timeline task view of the Health Care Assistant.

These two views offer different possibilities to the user, namely the possibility to get an overall view of the clinical process with the calendar view and to focus on a task at a time with the timeline view. The notifications mentioned throughout Section 3.4 can be seen as side messages, as shown in 8. By clicking on a task entry, it is possible to visualize task details such as remaining execution time and number of executions, task descriptions and so forth, as seen in Fig. 10.

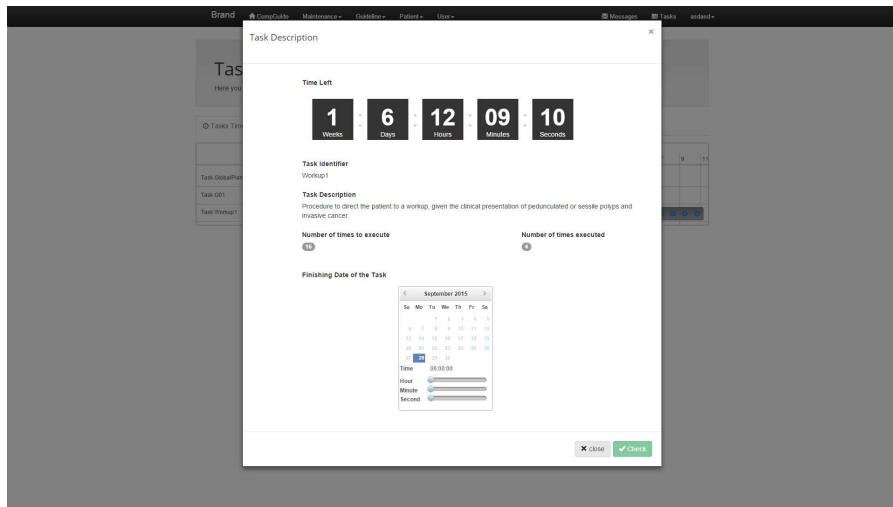


Fig. 10 Description of a clinical task in the Health Care Assistant.

4 Conclusions and Future Work

The work presented herein is an implementation example of the notion of guidelines as services, presented in [Novais et al, 2016], which takes advantage of the flexibility of the CompGuide system. The main contributions are a comprehensive temporal representation model and a web-based tool for the execution of CIGs. The tool builds an agenda of clinical tasks for the health care professional to follow and provides timely notifications of clinical events, while filtering the advice given to the health care professionals at a given time. The intention is to lessen the burden placed on them and help to keep their patients on the right track. Compared to current applications for the execution of CIGs, the one presented herein reflects a different view of guideline application and is endowed with functionalities that go beyond the simple display of clinical tasks. It is also a reminder system that may help the user to manage time and to ensure the enactment of procedures. Although there is a functional prototype, the tool is still in development. However, it represents a path to make CPGs more dynamic and improve their daily application.

As future work, it is necessary to evaluate the HCA tool by performing usability tests with health care professionals in order to infer about the usefulness of the developed application. A functionality that is currently being developed is the integration of clinical tasks with the calendar service used by the health care professional, thus enabling the visualization of CIG executions not only on the HCA, but also on their own calendar services, such as Google Calendar or Microsoft Calendar), with the other events of their daily practice which are outside the scope of guideline execution.

Acknowledgements This work has been supported by COMPETE: POCI-01-0145-FEDER-0070 43 and FCT Fundação para a Ciência e Tecnologia within the Project Scope UID/CEC/ 00319/2013. The work of Tiago Oliveira is supported by a FCT grant with the reference SFRH/BD/85291/ 2012.

References

- Adhikari NKJ, Beyene J, Sam J, Haynes RB (2005) Effects of Computerized Clinical Decision Support Systems on Practitioner Performance. *Journal of the American Medical Association* 293(10):1223–1238
- Anselma L, Terenziani P, Montani S, Bottrighi A (2006) Towards a Comprehensive Treatment of Repetitions, Periodicity and Temporal Constraints in Clinical Guidelines. *Artificial Intelligence in Medicine* 38(2):171–195, DOI 10.1016/j.artmed.2006.03.007
- Benson A, Bekaii-Saab T, Chan E, Chen YJ, Choti M, Cooper H, Engstrom P (2013) NCCN Clinical Practice Guideline in Oncology Colon Cancer. Tech. rep., National Comprehensive Cancer Network, URL http://www.nccn.org/professionals/physician_gls/f_guidelines.asp
- Berg D, Ram P, Glasgow J, Castro J (2004) SAGEDesktop: An Environment for Testing Clinical Practice Guidelines. *The 26th Annual Interna-*

- 1 tional Conference of the IEEE Engineering in Medicine and Biology Society
2 4:3217–3220, DOI 10.1109/IEMBS.2004.1403906
- 3
4 Boxwala Aa, Peleg M, Tu S, Ogunyemi O, Zeng QT, Wang D, Patel VL,
5 Greenes Ra, Shortliffe EH (2004) GLIF3: A Representation Format for
6 Sharable Computer-Interpretable Clinical Practice Guidelines. *Journal of*
7 *Biomedical Informatics* 37(3):147–61, DOI 10.1016/j.jbi.2004.04.002
- 8 Bright T, Wong A, (2012) Effect of Clinical Decision-Support Systems: A
9 Systematic Review. *Annals of Internal Medicine* 157(1):29–43, DOI 10.7326/
10 0003-4819-157-1-201207030-00450
- 11 Isern D, Moreno A (2008) Computer-based Execution of Clinical Guidelines: a
12 Review. *International Journal of Medical Informatics* 77(12):787–808, DOI
13 10.1016/j.ijmedinf.2008.05.010
- 14 Kaushal R, Shojania KG, Bates DW (2003) Effects of computerized physician
15 order entry and clinical decision support systems on medication safety: a
16 systematic review. *Archives of internal medicine* 163(12):1409–16, DOI 10.
17 1001/archinte.163.12.1409
- 18 McGuinness DL, Van Harmelen F (2004) OWL Web Ontology Language
19 Overview. URL <https://www.w3.org/TR/owl-features/>
- 20 Musen MA, Shahar Y, Shortliffe EH (2006) Clinical decision-support systems.
21 In: Shortliffe E, Cimino J (eds) *Biomedical Informatics, Health Informatics,*
22 *Springer New York*, pp 698–736, DOI 10.1007/0-387-36278-920
- 23 Novais P, Oliveira T, Neves J (2016) Moving towards a new paradigm
24 of creation, dissemination, and application of computer-interpretable
25 medical knowledge. *Progress in Artificial Intelligence* pp 1–7, DOI
26 10.1007/s13748-016-0084-2, URL [http://dx.doi.org/10.1007/
27 s13748-016-0084-2](http://dx.doi.org/10.1007/s13748-016-0084-2)
- 28
29 Oliveira T, Novais P, Neves J (2013) Representation of Clinical Practice Guide-
30 line Components in OWL. In: *Trends in Practical Applications of Agents*
31 *and Multiagent Systems SE - 10, Advances in Intelligent Systems and Com-*
32 *puting*, vol 221, Springer International Publishing, pp 77–85
- 33 Oliveira T, Leão P, Novais P, Neves J (2014) Webifying the Computerized Exe-
34 cution of Clinical Practice Guidelines. In: Bajo Perez J, Corchado Rodríguez
35 JM, Mathieu P, Campbell A, Ortega A, Adam E, Navarro EM, Ahrndt S,
36 Moreno MN, Julián V (eds) *Trends in Practical Applications of Heteroge-*
37 *neous Multi-Agent Systems. The PAAMS Collection SE - 18, Advances in*
38 *Intelligent Systems and Computing*, vol 293, Springer International Pub-
39 *lishing*, pp 149–156, DOI 10.1007/978-3-319-07476-4{-}18
- 40 Peleg M (2013) Computer-interpretable Clinical Guidelines: A Methodological
41 Review. *Journal of Biomedical Informatics* 46(4):744–63, DOI 10.1016/j.jbi.
42 2013.06.009
- 43 Samwald M, Fehre K, de Bruin J, Adlassnig KP (2012) The Arden Syntax
44 standard for clinical decision support: Experiences and directions. *Journal*
45 *of biomedical informatics* DOI 10.1016/j.jbi.2012.02.001
- 46 Shahar Y, Miksch S, Johnson P (1998) The Asgaard Project: A Task-specific
47 Framework for the Application and Critiquing of Time-oriented Clinical
48 Guidelines. *Artificial intelligence in Medicine* 14(1-2):29–51, DOI 10.1016/
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

S0933-3657(98)00015-3

- Silva A, Oliveira T, Novais P, Neves J (2015) Representing Temporal Patterns in Computer-Interpretable Clinical Guidelines. In: Schulz C, Liew D (eds) 2015 Imperial College Computing Student Workshop (ICCSW 2015), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, OpenAccess Series in Informatics (OASICS), vol 49, pp 62–69, DOI <http://dx.doi.org/10.4230/OASICS.ICCSW.2015.62>
- Sittig DF, Wright A, Osheroff Ja, Middleton B, Teich JM, Ash JS, Campbell E, Bates DW (2008) Grand challenges in clinical decision support. *Journal of biomedical informatics* 41(2):387–92, DOI 10.1016/j.jbi.2007.09.003
- Terenziani P, Montani S, Bottrighi A, Torchio M, Molino G, Correndo G (2004) The GLARE Approach to Clinical Guidelines: Main Features. *Studies in Health Technology and Informatics* 101(3):162–6, DOI 10.3233/978-1-60750-944-8-162
- Vollebregt A, ten Teije A, van Harmelen F, van der Lei J, Mosseveld M (1999) A study of PROforma, a development methodology for clinical procedures. *Artificial Intelligence in Medicine* 17(2):195–221, DOI 10.1016/S0933-3657(99)00016-0
- Wang D, Peleg M, Tu SW, Boxwala AA, Ogunyemi O, Zeng Q, Greenes RA, Patel VL, Shortliffe EH (2004) Design and implementation of the GLIF3 guideline execution engine. *Journal of Biomedical Informatics* 37(5):305–318, DOI 10.1016/j.jbi.2004.06.002