An approach for representing and managing medical exceptions in care pathways based on temporal hierarchical planning techniques*

Inmaculada Sánchez-Garzón[1], Juan Fernández-Olivares[1], and Luis Castillo[2]

[1] Department of Computer Science and A.I., University of Granada, Spain. {isanchez,faro}@decsai.ugr.es
[2] IActive Intelligent Solutions, Spain. {LCastillo}@iactive.es

Abstract. This work presents an approach for representing and managing medical exceptions that may occur during the execution of a patient-centered care pathway. Personalized care pathways are generated automatically by means of a knowledge-driven planning process over a temporal hierarchical task network (HTN), which encodes an evidence-based clinical guideline. The exceptional situations specified in this guideline as well as the recommendations for their management are represented by knowledge-based rules in the task network model. However these rules, which encode the exceptional flow of the guideline, are represented separately from the normal flow in order to not obscure the modelling. Moreover, we propose the use of medical concepts from a standard terminology (UMLS) for the formal representation of these rules. This fact promotes interoperability, knowledge sharing and precision aspects. Finally, a therapy planning system with capabilities for exception detection, analysis and adaptation has been developed. As a result, the proposal, which is evaluated with oncology care plans, seems to be an adequate exception recovery mechanism maintaining guideline adherence.

Keywords: Computer-interpretable clinical protocols, personalization and adaptation, exception handling, care pathway interoperability.

1 Introduction

Care pathways are defined as structured multidisciplinary care plans that detail the essential steps in the care of patients with a specific clinical problem [1]. Unlike clinical guidelines, care pathways do provide detailed guidance for developing and implementing each step of a care protocol based on clinical evidence. In this way, a care pathway specifies not only the timing and sequence of tasks (e.g., key events, clinical exams and assessments) to be performed on patients, but also describes operational and administrative actions of the organization considered in the care process, which in turn requires the involvement of

*This research work has been partially supported by the Andalusian Regional Ministry and the Spanish Ministry of Innovation under projects P08-TIC-3572 and TIN2008-06701-C03-02 respectively.

resources (e.g., clinical staff and medical equipment) [2]. The consideration of all these items (medical knowledge, patient data, availability of resources and temporal constraints) allows to design timed processes of patient-focused care that try to produce the best prescribed outcomes for an appropriate episode of care. Since the manual management of all these aspects is not trivial, IT support for the automated generation of care pathways is required. Moreover, this support should ensure that care pathways adhere to the recommendations specified in the clinical guideline. Regarding these issues, recent works [3, 4] have shown that the hierarchical planning paradigm (HTN) is an enabling technology to support processes in medical treatments as well as clinical decisions. This paradigm allows the representation of clinical guidelines as a temporal hierarchy of task network, like most of the CIGs formalisms [5]. In fact, [4] shows that temporal HTN is as expressive as the Asbru formalism. Clinical patient data and care organization details may also be represented. As a result, a knowledge-driven planning process over the task network model of the guideline together with the patient and organizational data, allows the generation of personalized care pathways automatically. Moreover, this planning process promotes the adherence between care pathways and the clinical guideline since it follows a deliberative reasoning process that integrates effective temporal constraints propagation techniques [6].

In addition to the IT support for generating care pathways automatically, the dynamism of healthcare domains forces that mechanisms for managing the medical exceptions that may occur during the execution of a care plan be also required. Influenced by the definitions proposed in [7, 8, 9, 10], we assume that a medical exception is any deviation from an ideal, normal care pathway and it involves a workflow adaptation for its resolution. The main sources of medical exceptions are changes in the clinical and social conditions of patients, new laboratory and toxicological findings and unexpected circumstances occurred inside the care organization. Moreover, the adaptation mechanism should be a knowledge-based adaptation process according to the recommendations specified in clinical guidelines as well as the experience of medical practitioners. Due to the difficulty on adjusting care pathways to exceptional situations taking all these requirements into account, techniques for representing and automatically handling exceptions during the enactment of a care plan are also required.

In this paper, we propose an approach for representing and managing medical exceptions that may occur during the execution of a care pathway using temporal HTN planning techniques [3, 4]. Every exceptional situation specified in the clinical guideline is associated with a declarative, knowledge-based rule. A rule is composed by four entities: (1) the medical exception detected, (2) the conditions under which this exception should be handled, (3) the sequence of steps to perform in order to adapt the care plan properly and (4) a descriptive field specifying the intention of the rule. The set of knowledge-based rules constitutes the *exceptional flow* of the clinical guideline and it is represented in the temporal task network model as a separate (from the normal flow) branch. Both the normal and the exceptional flows are represented by using a declarative and object-oriented language (EKDL[3]) that makes use of HTN as its underlying formalism [11, 6]. Moreover, in order to have an unambiguous representation of the knowledge-based rules, we propose to integrate medical concepts from a standard terminology. In this way, medical exceptions (e.g., diseases, symptoms, observations, findings) as well as the steps (e.g., care activities, laboratory procedures) and resources (e.g., healthcare professionals, medical equipment) required for adapting the care pathway are represented by UMLS [4] concepts.

Considering the exception management issue, we propose a therapy planning system whose operation mode is based on a continual planning approach [12]. The life cycle of this approach follows a tightly coupled control loop where tasks of *situation perception*, *reasoning* and *acting* are interleaved. As observed, this life cycle is also the underlying operation mode of a standard medical exception management system [13]. In this way, after the detection of a exceptional situation, the care pathway is adapted according to the knowledge-based rule in charge of its resolution. The idea is that the same knowledge-driven planning process used for the automated generation of personalized care pathways be also used for the automated adaptation of care plans. The main reason of this decision is that the resolution of a medical exception often requires appropriate reasoning and propagation techniques and this temporal planning process has shown to be effective in many knowledge-intensive domains [6, 3].

The experimental evaluation of the proposal here presented has been performed with care pathways of an oncology clinical guideline. As a result, some major contributions have been identified. Firstly, the separation between the normal flow definition and the flow representing the exceptional situations facilitates the readability and understanding of the guideline modelling. Secondly, the use of medical concepts from standard terminology allows to represent the knowledge-based rules in a precise way. As [14] remarks, this fact promotes semantic interoperability, knowledge shareability, patient safety and knowledge reuse. Moreover, the use of HTN planning techniques allows the automated adaptation of care pathways when a medical exception is detected. This planning process also ensures the adherence to the clinical guideline recommendations, which are encoded in the knowledge-based rules. Finally, the therapy planning system together with the rules provide capabilities to respond upon the occurrence of medical exceptions during the enactment of a care pathway, thus promoting aspects as adaptability and flexibility.

The paper is structured as follows. Section 2 describes the language and proposal for representing the knowledge-based rules. Section 3 presents the therapy planning system in charge of simulating, detecting and handling the medical exceptions. Then, in section 4, we describe the experimental evaluation that is focused on care pathways of the oncology area. Section 5 discusses related work. Finally, section 6 concludes the paper with some remarks and future work.

[3]Expert Knowledge Description Language is an evolution of the HTN-PDDL language, which was developed by our research group and now it is exploited by IActive Intelligent Solution `http://www.iactiveit.com/`.

[4]Unified Medical Language System. `http://www.nlm.nih.gov/research/umls/`

2 Exception representation by HTN planning techniques

The goal of this section is, firstly, to comment the main features of the HTN planning paradigm as well as the formalism used for the clinical guideline representation. Then, the approach for integrating standard medical terminology in the guideline modelling is presented. This assumption is the base for a precise representation of the knowledge-based rules, which will be described afterwards.

2.1 HTN planning paradigm and EKDL language.

Like most of the well-known CIGs languages [5], the HTN planning paradigm used in this work is an expressive approach for encoding both the clinical decisions and procedures found in a clinical guideline by a hierarchy of task network [3, 11]. This task network model describes how every *compound task* (e.g., apply a chemotherapy cycle) is decomposed into *atomic tasks* (e.g., drug administration actions) by applying a *decomposition method*. An atomic task is defined by a list of conditions, effects and parameters, which may represent the resources required for its execution or some additional qualifying data (e.g., drug administration mode). A method describes the set of conditions (e.g., the number of platelets, leukocytes and neutrophils in blood must be in the normal ranges) that must be satisfied by the clinical data in order for the method to be applicable by the planner. These clinical data are extracted from the VMR[5] of the care organization considered, which can also be represented in the HTN paradigm.

The formalism used for the definition of the task network model is the EKDL language, which covers the following desirable features from the knowledge representation standpoint. Firstly, it is a declarative language, which promotes both expressivity and inference aspects. It has support for temporal representation and reasoning since every compound and atomic tasks have start and end time points and complex temporal constraints can be specified over them [4, 6]. Moreover, conditions of methods and atomic actions are represented as logical expressions, which refer to an object-oriented data model that can be defined by using UML concepts. In addition, this language has a graphical notation[6] [15] that may promote the understandability of the clinical model for various kinds of users, as [7] remarks as a favourable feature. Finally, EKDL supports the definition of complex clauses (*macros*) that provide capabilities to infer new knowledge from the current clinical data. Regarding all these features, the EKDL language can be seen as an adequate formalism for clinical guideline representation.

2.2 Guideline modelling with standard clinical terminology.

Clinical guidelines detail instructions on how to perform diagnosis, therapy and follow-up, so these specifications constitute the general course of the careflow

[5] Virtual Medical Record

[6] EKMN is a notation inspired by BPMN, the current standard notation for process modelling.

(i.e., the normal flow). In addition, clinical guidelines often specify the steps to carry out in some *expected* exceptional situations. The approach here presented tries to use this information for the definition of knowledge-based rules that manage these exceptions. Since these rules are created from the explicit recommendations specified in clinical guidelines, the type of situations that our proposal can handle are *predictable* exceptions [8]. However, the high dynamism of healthcare domains may also produce *unpredictable* exceptions, that is to say, exceptions whose management is not specified in the clinical guideline. The goal here is to offer support to experts in order to guide them in the resolution of these exceptions. In these cases, we consider that the key is to share the knowledge encoded in other clinical guidelines and to offer a proper recommendation for facilitating the clinical decision task of experts. This fact requires mechanisms for knowledge retrieval as well as knowledge reuse. According to [14], the use of standard clinical terminologies facilitates such activities since they are composed by univocal concepts that are human-readable and computer-processable. So, our proposal is to integrate standard concepts in the modelling of the knowledge-based rules. The reason is to provide a precise representation that facilitates decision support after the occurrence of unpredictable exceptions.

Regarding the issue of standard terminologies, we propose the use of UMLS terms since this system covers a broad range of concepts in the healthcare area. The idea is that, as far as possible, every medical term involved in the guideline modelling be represented by its corresponding UMLS concept. In our approach, an UMLS concept is defined by its identifier (CUI), its name and its preferred semantic type. Moreover, in the case of concepts denoting a care activity, there is an extra field that links to the medical exceptions that can be detected by means of this activity. To clarify this proposal, a statement belonging to an oncology protocol will be analysed. As a result, the UMLS concepts extracted from it are represented in table 1.

Example 1. If after a blood laboratory test, the neutrophil count is below the normal range, growth factors must be applied to the patient. Growth factors require to administer (during 10 days and after the finalization of the current chemotherapy cycle) the drug Filgrastim with the recommended dose of 5 μgr/Kg. However, if the 7th day the total number of neutrophils is above 3000/ml, the administration of this drug is suspended. The administration mode of this drug can be subcutaneous or intravenous infusion (if the number of platelets is below 20000/ml).

CUI	Name	Semantic Type	Exceptions
C0009555	Complete Blood Count	Laboratory Procedure	C0580316
C0580316	Neutrophil count abnormal	Finding	
C0150270	Medication Management	Health Care Activity	
C0210630	Filgrastim	Pharmacologic Substance	
C0393123	Continuous infusion of chemotherapy	Therapeutic or Preventive Procedure	
C0812152	Medication administration: subcutaneous	Therapeutic or Preventive Procedure	

Table 1. UMLS concepts extracted from an oncology protocol. The care activity *Complete Blood Count* (C0009555) can detect the medical exception *Neutrophil count abnormal* (C0580316).

Once identified the relevant UMLS concepts, the next step is to define and model the clinical guideline by means of the HTN planning paradigm and the EKDL language. After that, elements used for the HTN-based guideline modelling will be tagged with its corresponding UMLS concept. Table 2 offers a summary of the correspondence between the main HTN elements and the typology of the UMLS concepts that they represent. For example, we can define an atomic task named *performRxNodes* that will be tagged with the UMLS concept whose identifier is C0024219. This means that this kind of task represents a *Diagnostic Procedure* whose formal name is *Lymphangiography*, i.e., a special X-ray of the lymph nodes. Applying this tagging process to every atomic task, parameter and medical exception; the specification of the clinical guideline (both the normal and exceptional flows) will be annotated with standard medical terminology. As commented before, this fact may result useful regarding the knowledge reuse aspect. For example, imagine that a patient is receiving a care plan for treating an oncology disease A. During this treatment, the patient suffers *acute stomatitis* and the oncology guideline A does not specify how to manage this disease. Maybe this exception is handled in an oncology guideline B and by means of the UMLS code of the exception (C0848336), the knowledge encoded in B can be easily processed and retrieved.

After the proposal for integrating standard terminology in the guideline modelling, the approach for handling medical exceptions will be described.

HTN Elements	UMLS Semantic Type	Example of UMLS Concept
Atomic task	Diagnostic Procedure	Lymphangiography (C0024219)
	Laboratory Procedure	Complete Blood Count (C0009555)
	Health Care Activity	Medication Management (C0150270)
	Therapeutic or Preventive Procedure	Platelet Transfusion (C0086818)
	Daily or Recreational Activity	Walking (C0080331)
	Educational Activity	Education and training activity (C0578156)
Resource parameters	Pharmacologic Substance	Prednisone (C0032952)
	Clinical Drug	Prednisone 2.5 MG Oral Tablet (C0690123)
	Medical Device	Brachytherapy device (C1299425)
	Professional or Occupational Group	Pediatric oncologist (C0279158)
Qualifying parameters	Body Part, Organ or Organ Component	Right lung (C0225706)
	Therapeutic or Preventive Procedure	Infusion of drug or medicament via intravenous route (C2317299)
Medical exceptions	Anatomical Abnormality	Bilateral small kidney (C0156246)
	Disease or Syndrome	acute stomatitis (C0848336)
	Finding	Platelet count abnormal (C0580317)
	Pathologic Function	Fibrosis (C0016059)
	Phenomenon or Process	Craniocerebral Trauma (C0018674)
	Sign or Symptom	Headache (C0018681)

Table 2. Correspondence between HTN elements and UMLS concepts.

2.3 Knowledge-based rules for handling medical exceptions.

The recommendations for managing medical exceptions are used for defining the declarative, knowledge-based rules that constitute the exceptional flow of the guideline. This flow is encoded as a compound task (called *Exceptional_Careflow*) in the task network model and it is decomposed in many branches as rules for handling exceptions are specified in the clinical guideline (see figure 1.b). Every knowledge-based rule is constituted by the following items:

- A reference to the UMLS concept(s) that formally represents the **medical exception(s)** that this rule tries to handle. Moreover, the use of UMLS identifiers (CUI) allows to specify exceptions univocally.
- A **list of conditions** that must be hold in the clinical data in order for the rule to be applicable. This entity is encoded as a decomposition method of the compound task *Exceptional_Careflow*.
- A **sequence of adaptation steps** for handling the exception properly. These steps may imply to add, delete, replace, adjust and postpone some activities in the care pathway. The way to encode these steps is by a **subtask network**. Moreover, as far as possible, every atomic task is tagged with the UMLS code of the care activity that it represents.
- A **description field** (or recommendation) that explains in a text-based format the intention of the knowledge-based rule. This item is relevant for facilitating the assistance to healthcare professionals and it is represented by an informative task at the beginning of the subtask network.

Regarding the situation specified in the example 1, two knowledge-based rules can be defined in the exceptional careflow. A more detailed description of these rules are represented in table 3.

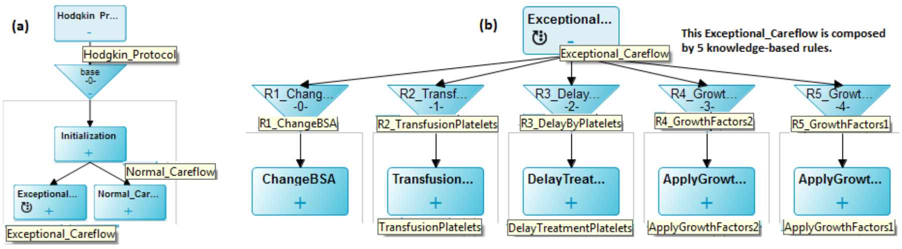

Fig. 1. Graphical representation of (a) a clinical guideline composed by its exceptional and normal careflows and (b) the knowledge-based rules composing an exceptional careflow.

Knowledge-based rule for applying growth factors (R5_GrowthFactors1)	
Exception	C0580316 (Neutrophil count abnormal).
Conditions	Current treatment phase of the patient is a chemotherapy cycle.
Adaptation Steps	(1) Recommendation. (2) Synchronization with the end date of the current chemotherapy cycle. (3) Updating current treatment phase with *GrowthFactors*. (4) Estimation of the drug administration mode (intravenous or subcutaneous) according to the level of platelets. (5) Generation 6 administration actions of filgrastim (once per day) 5 μgr/Kg. (6) Generation neutrophil test (7th day).

Knowledge-based rule for applying growth factors (R4_GrowthFactors2)	
Exception	C0580316 (Neutrophil count abnormal).
Conditions	Current treatment phase of the patient is *GrowthFactors* AND neutrophil count below 3000/ml.
Adaptation Steps	(1) Recommendation. (2) Estimation of the administration mode (intravenous or subcutaneous) according to the level of platelets. (3) Generation 4 administration actions of filgrastim (once per day) 5 μgr/Kg. (4) Synchronization with the start of the next chemotherapy cycle (temporal interval between cycles is 28 days).

Table 3. Knowledge-based rules for applying growth factors.

In our previous work [16], the exceptional flow is embedded in the normal flow definition. This approach obscures the guideline modelling since compound tasks in charge of handling the medical exceptions should be included in all those parts of the workflow in which the exceptions may occur, thus making model readability difficult. For this reason, the proposal of this paper is to represent both the normal and exceptional flows in the same task network model, but in separate branches (see figure 1.a). To clarify this approach, an intuitive example will be provided. The resolution of the exceptional situation of the example 1 is based on applying growth factors at the end of the current chemotherapy cycle. Considering the modelling in the *embedded* version, after every compound task representing a chemotherapy cycle (i.e., *ChemoCycle* task), a compound task in charge of checking and applying growth factors (when needed) must be added. As a result, many GrowthFactors tasks must be included in the guideline modelling as ChemoCycle tasks are defined (see figure 2.a). The same situation happens for the rest of exceptions, thus obscuring the modelling. In contrast, the proposal here presented simplifies the guideline modelling since only a GrowthFactors task is needed in order to handle exceptions related to abnormal count of neutrophils during the enactment of care pathways (see figure 2.b).

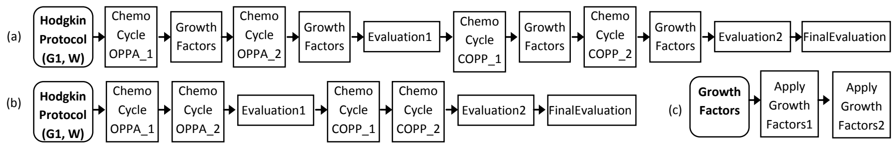

Fig. 2. Normal flows representing oncology treatments for a patient whose profile is woman (gender) and G1 (staging group). (a) Embedded version. (b) Separate version. (c) Decomposition of the GrowthFactors task. In the embedded version, this task is defined in the exceptional careflow.

As shown, the separation between the normal and exceptional flows facilitates the readability of the model. Moreover, regarding the typology of exceptions defined in [8], the exception handling approach here presented is able to detect predictable, synchronous and asynchronous exceptions. For this last purpose, the exception handler requires that the clinical data of the patient be always updated with the last steps performed. This is the way to know what is the exact patient state in order to synchronize the pending activities properly.

At this point, both the normal flow of the clinical guideline and its exceptional situations are represented by means of a temporal hierarchical planning approach. Covered the exception representation issue, the next goal is to face the exception management aspect, which will be presented in the next section.

3 Managing exceptions by a continual planning approach

In general terms, exception management during the enactment of a care pathway implies three steps: (1) to **perceive** the situation (in order to detect exceptional situations); (2) to **reason** from the current clinical situation perceived (i.e., to decide if it is possible to continue with the execution of the care plan or it must be adapted); and (3) to **act**, thus executing every activity of the care pathway step-by-step. Since this life cycle is the underlying operation mode of the continual planning approach, we propose a continual therapy planning system [12, 16] as architecture for managing exceptions during the execution of patient-centered care pathways. As our previous work [16] remarks, this architecture promotes the adaptive execution of care pathways. Moreover, a VMR can be used as an intermediate gateway integrating patient clinical data. Figure3 offers an overview of the therapy planning system developed, whose main components are:

Care pathway execution. This component is in charge of keeping the execution trace of care pathways. Moreover, after the execution of every care activity, the content of the VMR is updated properly.

Care pathway monitoring. The main mission of this component is to supervise that everything is going as planned, that is to say, to compare the expected health conditions of the patient with the real health conditions, which are extracted from the VMR. By means of this checking, medical exceptions may be detected autonomously.

Medical exception handling. Once the monitoring component detects any deviation in the patient clinical data, an exception handling step is carried out. The idea is to create a new record of exception in the VMR, which will be linked to the patient. A medical exception instance is defined by the following fields: (1) the UMLS code of the exception; (2) the context information in which the exception occurs (treatment phase, start and end dates); (3) the current status of the exception (*detected, pending, solved*).

Care pathway adaptation. The resolution of a medical exception often requires effective reasoning and propagation techniques in order to adapt the care pathway properly. For this reason, the idea is that the same knowledge-driven planning process proposed for the automated generation of care pathways, be also used for the automated adaptation of the care plan when a medical exception is detected. This means that a new replanning episode is carried out. However, in this case, the content of the VMR has been updated with a new record of medical exception, which has not been solved yet (i.e., its status is *detected* or *pending*). By means of this new record, the planning process is aware that the task representing the exceptional careflow must be applied for solving the exception. Then, the pending part of the treatment is generated by the task defining the normal flow. As a result, the care pathway is adapted according to the steps specified in the corresponding knowledge-based rule. After that, the adapted care plan resumes its execution.

The current setting of this architecture does not provide support for human interaction. However, its modular composition allows the integration of a component in charge of this requirement easily. In fact, in our previous work [11], care plans are integrated in BPM runtime engines. These tools are often composed by visual consoles, which promote human interaction. So, this proposal can be used for supporting interactive execution of care plans. Anyway, the care pathway adaptation process is able to handle both exceptions that have been raised by the monitoring component (autonomously) or by physicians (manually).

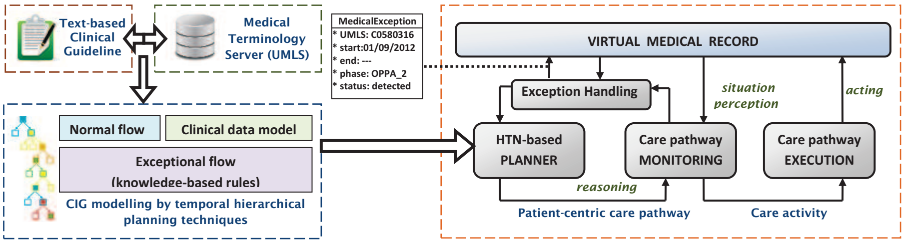

Fig. 3. Architecture of the continual therapy planning system.

4 Experimental Evaluation

For the experimental evaluation of this work, we have selected a clinical protocol of the oncology area[7]. This protocol constitutes long running care plans with complex temporal constraints and fixed patterns of drug administration and radiation sessions, whose dose and intensity depend on the patient profile. Moreover, it details the steps to perform in some expected exceptional situations.

In the first place, the medical entities required for representing the guideline in a precise way have been identified, giving as a result 63 UMLS concepts (see figure 4.a). Then, both the normal flow and the exceptional situations specified in the protocol have been modelled with the EKDL formalism, as explained in section 2. The normal flow encodes the care treatment of six different patient profiles, according to the staging group (G1, G2 or G3) and the gender (M or W) of the patient. The exceptional flow is composed by 10 knowledge-based rules (see table 4) that manage 13 out of 15 predictable exceptions identified (see figure 4.b). The reason is that the oncology guideline does not specify how to act in the presence of the exceptions *acute stomatitis* and *Graft vs Host Reaction*, which are defined as *non-specified predictable* exceptions.

Secondly, we have defined 6 different patients (one of each profile) and we have generated a personalized care pathway for each of them. These care plans are integrated in the continual therapy planning system presented in section 3 to be executed step-by-step. In order to force the appearance of medical exceptions, as if they occurred in a real healthcare environment, a simulation module has been included in the system (between the execution component and the VMR). This component is in charge of generating medical exceptions in a probabilistic way, but taking into account the set of possible exceptions that may occur at each step. As explained in section 2.2, if the *complete blood count* procedure is linked to the medical exception *abnormal levels of neutrophils*, the simulation module may generate this *possible* exception after the execution of this action.

As a result, some major contributions may be identified. Firstly, by the separation of the normal and exceptional flows, the guideline modelling has been

[7]Protocol for diagnosis and treatment the Hodgkin's disease in childhood and adolescents `https://decsai.ugr.es/~isanchez/EH-SEOP.003.pdf` (in Spanish) defined by the Spanish society of paediatric oncology (SEOP: `http://www.sehop.org/`).

simplified. In the separate version, the model is composed by 63 compound tasks (i.e., 11 tasks of the exceptional flow, 18 chemotherapy cycles Cx, 12 radiotherapy cycles Rx, 20 evaluation cycles Ex, 2 other tasks). However, if we consider the worst case of the embedded version in which tasks for handling the exceptions must be included in all those parts of the workflow definition in which they may occur, the total number of tasks increases to 278. Secondly, the proposal seems to be an adequate mechanism to promote the adaptive execution of care pathways since all the predictable exceptions can be autonomously detected and handled according to the specification of the clinical guideline. However, the issue of non-specified predictable exceptions as well as unpredictable exceptions will be studied in our future work. For this purpose, we consider that the use of standard medical terminology will be useful for knowledge reuse.

(a)

Type of UMLS concept	Total
Finding and Diseases	15
Laboratory procedures	7
Diagnostic procedures	12
Therapeutic procedures	7
Healthcare professionals	5
Medical equipment	1
Drugs	12
Administration modes	4
	63

(b)

CUI	Name	CUI	Name
C0015967	Fever	C0580317	Platelet count abnormal
C0027947	Neutropenia	C0580531	White blood cell count abnormal
C0021311	Infection	C0580316	Neutrophil count abnormal
C0043094	Weight Gain	C0586989	Varicella-zoster virus infection
C0424644	Has grown in height	C1968787	Peripheral motor neuropathy, severe
C0043096	Body Weight decreased	C1968788	Peripheral sensory neuropathy, severe
C0015230	Exanthema	C0018134	Graft vs Host Reaction
C0848336	acute stomatitis	**15 Medical Exceptions**	

Fig. 4. (a) UMLS concepts extracted from the Hodgkin Protocol. (b) Medical exceptions.

5 Related Work

Regarding the issue of representing clinical guidelines in a computer interpretable format, a broad range of formal languages has been developed [5, 17]. However, the runtime engines of these formalisms do not include support for the automated generation of patient-tailored care pathways, even less considering temporal and resource constraints. In addition, few languages offer support for the integration of standard terminology at the knowledge representation phase. While GUIDE language represents each task with a single SNOMED[8] code, GLIF3 is composed by a layer that provides access to medical knowledge databases.

Other key topic in the formulation of clinical guidelines is the representation of medical exceptions. In order to address this topic, new languages have been developed, as CIGDec [18]. The key of this declarative language is to specify what to do and leave it up to the user to decide how to work depending on the case. However, this formalism is not recommended for modelling strict and large processes, as is the case of oncology processes. [8] offers a complete conceptual model of exceptions using the Object Process Methodology (OPM), but this approach neither supports the specification of recovery mechanisms to handle the exceptions nor the resumption to normal execution mode of the care pathway.

The exception management issue is covered in [19], where the representation of exceptions is based on the PROforma formalism, which is an executable language. This work proposes a mechanism for handling exceptions during the

[8] `http://www.nlm.nih.gov/research/umls/Snomed/snomed_main.html`

execution of a CIG. However, our approach is based on managing exceptions that occur during the execution of a care pathway, which is automatically generated from a CIG by means of HTN planning techniques. For this reason, our handling process ensures that the adapted care plan is also patient-centric.

Considering the area of workflow systems, many efforts have been developed in order to both represent and handle exceptions that can occur during the enactment of a workflow. [10] proposes a taxonomy of exceptions and [7] presents a brief review related to exception management in medical workflow systems. As this work remarks, most proposals in this area are based on stand-alone approaches since they try to separate the exception management from the workflow process definition. In addition, the more common way to implement this approaches is by ECA (Event-Condition-Action) rules [9, 20, 21, 22] or by semantic rules [23]. Very relevant in this area is the ADEPT project [24], which describes the key requirements and challenges to be considered in healthcare processes. Regarding the most related works, [21] and [22] present some limitations in modelling complex time constructions. This fact is supported by our HTN planning paradigm [6, 4]. Moreover, these related works are based on handling predictable exceptions since the ECA rules are created from the detailed instructions specified in clinical trials. However, as presented in this work, the management of unpredictable exceptions is also relevant in healthcare domains.

Rule	**Medical Exceptions**	**Description**	**Adaptation Steps**					**Phase**		
			add	del	delay	synch	adj	Cx	Rx	Ex
R1	C0424644 C0043096 C0043094	Calculate the body mass index of the patient and adjust the dose of the pending drug administration activities.					✓	✓	✓	✓
R2	C0580317	Delay the treatment (if first episode) and the current phase is Cx or Rx.			✓			✓	✓	
R3	C0580317	Perform platelet transfusion (if more episodes) and resume treatment next day.	✓			✓		✓	✓	✓
R4	C0580531	Delay the treatment (if first episode) and the current phase is Cx or Rx.			✓			✓	✓	
R5	C0580531	Perform white cells transfusion (if more episodes) and resume treatment next day.	✓			✓		✓	✓	✓
R6	C0580316 C0027947	Growth factors 1: administer filgrastim every 24h. during 6days (5 μgr/kg) and perform a neutrophil test the 7th day.	✓			✓		✓		
R7	C0580316 C0027947	Growth factors 2. Administer filgrastim every 24h. during 4days (5 μgr/kg).	✓			✓		✓		
R8	C1968787 C1968788	Suspend one administration of the drug which causes the disease and resume the next ones at 50 % of the recommended dose (in current chemotherapy cycle).		✓			✓	✓		
R9	C0586989 C0015230	Administer acyclovir (intravenous infusion 1h.) every 8h. during 7days (5mg/kg) and perform test varicella.	✓			✓		✓	✓	✓
R10	C0015967 C0021311	Administer one tablet of cotrimoxazole (oral) every 12h. during the current phase. If tumour location is spleen and patient will receive RT: Perform antipneumococcal and meningococcal C vaccines (before RT). Administer Penicillin V (oral) every 12h with dose 125 mg (age$\leq$5) or 250 mg (age>5) after RT.	✓			✓	✓	✓	✓	✓

Table 4. Knowledge-based rules extracted from the Hodgkin Protocol. For every rule, the exceptions that handles, the adaptation steps (add, delete, delay, synchronise, adjust) and the threatment phase in which the exception may occur (chemotherapy, radiohterapy, evaluation) are specified.

Taking the HTN planning area into account, several works try to adapt plans that fail during an execution episode [25, 26]. The key of these proposals is to find another way to carry out the action that has failed. However, these proposals are far from being used in healthcare configurations where the adaptation must be based on the knowledge specified in the clinical protocol and it is necessary to consider all the previous steps performed to the patient.

Comparing the proposal here presented with our previous works [3, 11, 16], some contributions can be identified. On the one hand, the key of [3] is to prove that HTN paradigm is an enabling technology for both representing clinical guidelines and generating personalized care pathways automatically. This work does not represent or manage any exceptional situation. On the other hand, [11] presents a knowledge engineering suite[9] based on the HTN technology that facilitates the clinical guideline modelling. Although this work presents techniques for executing and monitoring care plans, the exception handling approach is still at a conceptual design level. Finally, in [16], neither the guideline is represented by means of standard clinical concepts nor there is separation between the exceptional and the normal flows. So, the guideline representation is obscured since the exceptional knowledge is embedded in the normal flow. The main goal of [16] is to present a continual planning architecture for promoting the adaptive execution of care pathways, so the representation issue is not a primordial aspect.

6 Conclusions and Future Work.

This work proposes an approach for representing and managing medical exceptions that may occur during the execution of a patient-centered care pathway. This approach is based on temporal hierarchical planning techniques and it offers some advantages. Firstly, the HTN paradigm is an enabling technology for representing and managing medical exceptions, thus ensuring the adherence between care pathways and clinical guideline recommendations. Secondly, the separation between the normal and exceptional flows allows not to obscure the guideline modelling. Moreover, HTN paradigm also allows the integration of medical concepts from standard terminology, thus promoting the guideline representation in a precise form. Finally, the experimental evaluation shows that a complete and standard representation of clinical guidelines together with a therapy planning system with capabilities for exception detection and exception handling, may result in an effective approach to support the adaptability of clinical pathways.

Regarding the future work, we mention some ideas. Firstly, since the integration of UMLS concepts in the guideline modelling is made manually, mechanisms for importing concepts from a terminology server automatically may be very advantageous. Secondly, we will try to manage unpredictable exceptions [8]. For this reason, our first goal is to reuse the knowledge encoded in rules of other clinical guidelines for offering assistance to healthcare professionals. In addition, the use of standard terminology promotes this requirement. Finally, since the

[9] IActive Knowledge Studio. `http://www.iactiveit.com/technology/studio-expert-knowledge-modeling-tool/`

participation of experts is key in clinical contexts, a component for promoting human interaction will be integrated in the therapy planning system proposed.

References

[1] H. Campbell, R. Hotchkiss, N. Bradshaw, and M. Porteous. Integrated care pathways. *British Medical Journal*, 316(7125):133–137, 1998.

[2] G. van Merodea, S. Groothuisb, and A. Hasmanb. Enterprise resource planning for hospitals. review. *Int. J. of Medical Informatics*, 73(6):493–501, 2004.

[3] J. Fdez-Olivares, L. Castillo, J. A. Cózar, and O. García-Pérez. Supporting clinical processes and decisions by hierarchical planning and scheduling. *Comp. Intelligence*, 27:103–122, 2011.

[4] A. González-Ferrer, A. ten Teije, J. Fdez-Olivares, and K. Milian. Automated generation of patient-tailored care pathways by translating computer-interpretable guidelines into hierarchical task networks. *Artificial Intelligence In Medicine (to appear)*, 2012.

[5] M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R. A. Greenes, R. Hall, P. D. Johnson, N. Jones, A. Kumar, S. Miksch, S. Quaglini, A. Seyfang, E. H. Shortliffe, and M. Stefanelli. Comparing computer-interpretable guideline models: A case-study approach. *J. Am. Med. Inform. Assoc*, 10(1):5268, 2003.

[6] L. Castillo, J. Fdez-Olivares, O. García-Pérez, and F. Palao. Efficiently handling temporal knowledge in an HTN planner. In *Proc. of ICAPS06*, pages 63–72, 2006.

[7] M. Han, T. Thiery, and X. Song. Managing exceptions in the medical workflow systems. In *Proc. of the 28th International Conference on Software Engineering*, pages 741–750, 2006.

[8] M. Peleg, J. Somekh, and D. Dori. A methodology for eliciting and modeling exceptions. *Biomedical Informatics Journal*, 42:736–747, 2009.

[9] M. Klein and C. Dellarocas. A knowledge-based approach to handling exceptions in workflow systems. *Computer Supported Cooperative Work*, 9:399–412, 2000.

[10] F. Casati and G. Cugola. Error handling in process support systems. *LNCS. Advances in Exception Handling Techniques*, 2022:251–270, 2001.

[11] J. Fdez-Olivares, I. Sánchez-Garzón, A. González-Ferrer, J.A. Cózar, A. Fdez-Teijeiro, M.R. Cabello, and L. Castillo. Task network based modeling, dynamic generation and adaptive execution of patient-tailored treatment plans based on smart process management technologies. *LNCS. Knowledge Representation for Health-Care*, 6924:37–50, 2012.

[12] C. Guzmán, V. Alcázar, D. Prior, E. Onaindia, D. Borrajo, J. Fdez-Olivares, and E. Quintero. PELEA: a domain-independen architecture for planning, execution and learning. *Scheduling and Planning Applications woRKshop. ICAPS conf. (to appear)*, 2012.

[13] S. Miksch and A. Seyfang. Continual planning with time-oriented, skeletal plans. In *Proc. of ECAI*, 2000.

[14] Institute of Medicine. *Patient Safety: Achieving a New Standard for Care*. Washington (D.C.): National Academy Press, 2004.

[15] F. Palao, J. Fdez-Olivares, L. Castillo, and O. García. An extended HTN knowledge representation based on a graphical notation. *Workshop on Knowledge Engineering for Planning and Scheduling. ICAPS conf.*, pages 126–135.

[16] G. Milla-Millán, J. Fdez-Olivares, I. Sánchez-Garzón, D. Prior, and L. Castillo. Knowledge-driven adaptive execution of care pathways based on continuous planning techniques. *Workshop ProHealth-KR4HC. BPM conf. To appear*, 2012.

[17] P. de Clercq, J. Blom, H. Korsten, and A. Hasman. Approaches for creating computer-interpretable guidelines that facilitate decision support. *Artif Intell Med*, 31(1):1–27, 2004.

[18] N. Mulyar, M. Pesic, W. van der Aalst, and M. , Peleg. Declarative and procedural approaches for modelling clinical guidelines: Addressing flexibility issues. *Business Process Management Workshops. Lecture Notes in Computer Science*, 4928:335–346, 2008.

[19] A. Grando, M. Peleg, and D. Glasspool. A goal-oriented framework for specifying clinical guidelines and handling medical errors. *Biomedical Informatics Journal*, 43:287–299, 2010.

[20] Z. Luo, Kochut K. Sheth, A., and J. Miller. Exception handling in workflow systems. *Applied Intelligence*, 13:125–147, 2000.

[21] U. Greiner, R. Müller, E. Rahm, J. Ramsch, B. Heller, and M. Loeffler. AdaptFlow: Protocol-based medical treatment using adaptative workflows. *Methods of Information in Medicine*, 44:80–88, 2005.

[22] R. Müller, U. Greiner, and E. Rahm. AgentWork: a workflow system supporting rule-based workflow adaptation. *Data and Knowledge Eng.*, 51:223–256, 2004.

[23] D.A. Alexandrou, I.E. Skitsas, and G.N. Mentzas. A holistic environment for the design and execution of self-adaptive clinical pathways. *Information Technology and Applications in Biomedicine Conference (ITAB)*, pages 1–5, 2009.

[24] P. Dadam and M. Reichert. The ADEPT project: a decade of research and development for robust and flexible process support. Challenges and achievements. *Computer Science. Research and Development.*, 23(2):81–97, 2009.

[25] N.F. Ayan, U. Kuter, F. Yaman, and R.P. Goldman. HOTRiDE: Hierarchical ordered task replanning in dynamic enviroments. In *Proc. of ICAPS*, 2007.

[26] I. Warfield, C. Hogg, S. Lee-Urban, and H. Muñoz Avila. Adaptation of hierarchical task network plans. In *Proc. of FLAIRS*, pages 429–434, 2007.