Personalized conciliation of clinical guidelines for comorbid patients through multi-agent planning☆

Juan Fdez-Olivares[a,*], Eva Onaindia[b], Luis Castillo[a], Jaume Jordán[b], Juan Cózar, MD[c]

[a]*Department of Computer Science and AI, Universidad de Granada, Spain*
[b]*Department of Computer Science, Universitat Politècnica de València, Spain*
[c] *Hospital San Agustín, Linares, Spain.*

Abstract

The conciliation of multiple single-disease guidelines for comorbid patients entails solving potential clinical interactions, discovering synergies in the diagnosis and the recommendations, and managing clinical equipoise situations. Personalized conciliation of multiple guidelines considering additionally patient preferences brings some further difficulties. Recently, several works have explored distinct techniques to come up with an automated process for the conciliation of clinical guidelines for comorbid patients but very little attention has been put in integrating the patient preferences into this process.

In this work, a Multi-Agent Planning (MAP) framework that extends previous work on single-disease temporal Hierarchical Task Networks (HTN) is proposed for the automated conciliation of clinical guidelines with patient-centered preferences. Each agent encapsulates a single-disease Computer Interpretable Guideline (CIG) formalized as an HTN domain and conciliates the decision procedures that encode the clinical recommendations of its CIG with the decision procedures of the other agents' CIGs. During conciliation, drug-related interactions, scheduling constraints as well as redundant actions and multiple support interactions are solved by an automated planning process. Moreover, the simultaneous application of the patient preferences in multiple diseases may potentially bring about contradictory clinical decisions and more interactions. As a final step, the most adequate personalized treatment plan according to the patient preferences is selected by a Multi-Criteria Decision Making (MCDM) process. The MAP approach is tested on a case study that builds upon a simplified representation of two real clinical guidelines for Diabetes Mellitus and Arterial Hypertension.

Keywords:
Clinical Guidelines, Comorbidities, Concilitation, Patient Preferences, HTN Planning, Multi-Agent Planning

1. Introduction

Clinical Practice Guidelines (CPGs) are statements that include recommendations to assist clinical decision making in patient care. Recent clinical research unveils a widening gap between the reality of healthcare of comorbid patients and the practical clinical recommendations driven by CPGs [1]. A number of distinct approaches for automatically conciliating CPGs of different diseases that need to be simultaneously applied in a patient have emerged

☆This work has been partially supported by Spanish Government Projects MINECO TIN2014-55637-C2-2-R and TIN2015-71618-R.
*Corresponding author
Email addresses: faro@decsai.ugr.es (Juan Fdez-Olivares), onaindia@dsic.upv.es (Eva Onaindia), L.Castillo@decsai.ugr.es (Luis Castillo), jjordan@dsic.upv.es (Jaume Jordán), jcozarolmo@hotmail.com (Juan Cózar, MD)

lately, where the common mainstream relies in representing each single-disease Computer Interpretable Guideline (CIG) in a suitable formalism, finding merging points between the CIGs and tackling the potential arising interactions.

Approaches to comorbidity vary depending on various factors such as the knowledge-engineering tools to formalize the individual CPGs, the type of representational formalism and the reasoning process. Despite the great variety of models and usage of diverse technology in comorbidity, little effort has been devoted to handle comorbid patients preferences. Most likely, the reason stems from the additional complexity of dealing with the contradictions and adverse interactions that may arise when conciliating clinical decisions that comply with the patient preferences for different diseases.

The inherently complexity of adding patient preferences in comorbidity settings depends to a greater or lesser extent on the reasoning process of the comorbid approach. Approaches that merge multiple treatments for a patient, managing actions and enforcing constraints encoded in each treatment [2], are less flexible because the appearance of an adverse interaction when a particular merge criteria is applied on the patient compels to change one of the treatments. On the other hand, proposals that adopt the trend of conciliating CIGs instead of treatments cannot easily accommodate the patient preferences during conciliation since CIGs are designed for diseases, not for patients [3]. These approaches are specifically aimed to enable abstract or patient-independent reasoning in order to find possible interactions between CIGs before execution [4].

The generation of patient-oriented CIGs is gaining much attention as it enables mobile decision support [5, 6]. Recent developments in the MobiGuide project apply CIGs to provide real-time patient-specific and personalized recommendations by matching CIG knowledge with a highly-adaptive patient model, thus making CIGs patient-centered and enabling their personalization [7]. Despite these advancements in accommodating personal preferences in individual CIGs, conciliating several CIGs in a comorbidity setting also involves conciliating guideline-specific preferences, which may render new adverse interactions (e.g., two different drugs cannot be simultaneously taken at the patient preferable scheduled time). This opens up a new spectrum of *preference interactions* as a result of the simultaneous application of the patient preferences to all the diseases, which may also have a significant impact from a medical standpoint. With these antecedents in mind, it is understandable that efforts in comorbidity investigation have been focused on the detection of disease interactions rather than on conciliating patient preferences.

In this paper, we present a framework for conciliating clinical guidelines that overcomes the difficulties of previous approaches in managing patients preferences. Unlike the widely extended patient-centered treatments, our work represents a step ahead towards *patient-tailored conciliated clinical guidelines* in comorbidity. We propose a dynamic conciliation of CIGs that not only handles relevant medical interactions addressed by other approaches but also studies the potential harmful interactions that arise when the patient preferences are applied to more than one recommendation (preference interactions).

Beyond the medical state of the patient, which is the primary data source for medical decisions, patient preferences relate to a number of personal choices that denote the extent to which given health states are desirable [8, 9]. In this work, we handle two types of personal preferences:

- Qualitative preferences are personal choices of the patient that reflect common patterns of everyday life, habits or lifestyle and are used to involve patients in the development of the CPG
- Quantitative preferences are personal choices of the patient (or institution) that reflect the extent to which some treatment guidelines are less or more preferable than others

Our approach accounts for patient preferences at different stages of the health-care process, thus providing the most adequate multi-disease personalized and conciliated treatment:

1) qualitative preferences are represented in a predicate-based formalism and activate alternative decision procedures in each individual disease CIG, which is formally defined by the knowledge engineer as a planning domain (e.g., reasoning procedures differ according to the tolerance of the patient to common side effects of a medication or to the prescribed medication the patient is accustomed to). Preferences embodied in the individual guidelines are a helpful mechanism when designing a treatment.

2) conciliation of the most relevant interactions managed by clinicians [10] during the treatment elaboration, including those that stem from conciliating simultaneously guideline-specific personal preferences (preferences interactions).

3) resolving clinical equipoise situations, i.e., selecting one treatment from many possible clinically valid treatments, through the use of quantitative preferences.

Our approach draws upon the use of multi-agent techniques and multi-agent planning (MAP) in particular, in which a planning agent is regarded as encapsulating the expertise and skills of a specialist or team of specialists in a particular disease [11]. The adoption of MAP technology is very suitable for a comorbidity setting for several reasons: (a) there is no need to mix knowledge of different guidelines in a single representation; each agent individually encapsulates the CPG of a particular disease; (b) ease of including as many diseases as the patient suffers; and (c) support of existing and well-studied mechanisms of agent coordination and conflict resolution. Additionally, the strength of deliberative temporal Hierarchical Task Network (HTN) planning reveals as a solid enabling technology to represent CIGs and generate personalized single-disease treatments [12]. Thus, an agent-based HTN knowledge base is used to represent the clinical expertise of the agents (decision procedures) as well as to account for the patient preferences.

Each agent maintains an single-disease CIG represented as an HTN domain. Coordination among agents is aimed at conciliating the agents' local CIGs, which is performed in two steps. The first step, *domain merging*, requires an agent to filter out the decision procedures that are not compliant with the patient preferences and then combine the resulting sub-domain with the sub-domains of the other agents. Therefore, each agent ends up with its own unified merged CIG, resulting from the combination of the agents' decision procedures while respecting and conciliating the contents of each CIG. Moreover, if two decision procedures are found to be in conflict due to a preference interaction within one agent, the agent will request the other one an alternative procedure to resolve the interaction. At the second step, *the planning process*, every agent generates a personalized plan with its embedded local planner, using its conciliated CIG and considering the patient comorbid medical state and preferences. Computing a plan entails solving the adverse interactions that appear when the combined decision procedures are applied to the particular patient. If a planning agent finds a plan, meaning the problem is solvable, a local conciliated plan (treatment) has been encountered for the patient. Quantitative preferences are then applied to all the found local conciliated treatments in order to select the final multi-disease treatment plan.

A distinctive feature of our approach is the comprehensive temporal planning process applied by the agents. This encompasses reasoning about both the temporally annotated patient data and explicit temporal constraints represented in the HTN knowledge base, as well as handling the implicit temporal constraints derived from the hierarchical task decomposition at different levels of abstraction. Health conditions of the patient over time or prescribed treatments

dates are easily described with the language used to formalize CIGs as temporal HTN domains. Moreover, implicit temporal consistency checking is used to determine the temporal instantiation of actions or the fulfillment of resource availability restrictions. This powerful temporal reasoning machinery enables as well to infer the temporal horizon of a treatment plan in the form of a patient follow-up date.

From a technical standpoint, we propose a novel agent-based planning method to tackle the problem of personalizing treatments for comorbid patients accounting for personal and treatment choices of the patient. Our approach highlights the distributed representation of clinical knowledge and the coordination among agents as key elements to conciliate clinical guidelines, enabling specialists to exchange their knowledge to come up with a patient-centered treatment plan. Moreover, it also analyzes equipoise situations when several clinically valid alternatives exist for the same patient.

The outcome of our tool is an automatically built conciliated plan which is then shown to the physician. While the physician does not interactively intervene during the construction of the plan, the decisions made by the HTN planner are guided by human knowledge encoded by knowledge engineers in collaboration with physicians. Since the reasons leading to the output plan are internally recorded, they would be readily retrievable and accessible to the physician in the form HTN-based explanations through a suitable human-computer interface. This way, automated processes that follow a human-centered knowledge modeling like ours enable the generation of transparent and explainable decisions.

We must note that in this paper we do not address the issue of plan monitoring, plan execution (treatment execution) or the interaction with the patient, beyond the interaction required to capture the patient preferences. All in all, our proposal is an *enabling technology*; i.e., a methodology that, in combination with other associated technologies like user interaction, plan execution or computer-physician interaction, can provide the means to develop clinical decision systems. More details on how the tool is extensible to account for these interactive technologies will be exposed in the section devoted to discuss some open issues.

In the remainder of this paper, we will see the benefits of exploiting MAP technology for our comorbidity approach. This paper is organized as follows. In the next section, we summarize the principal approaches to comorbidity. Section 3 introduces the principal concepts of our HTN-based knowledge representation formalism. Section 4 overviews our patient preference classification. Section 5 outlines the architecture for the proposed MAP system and it explains the overall workflow. Section 6 presents the two steps of our MAP proposal, the domain merging and planning processes. The following section is devoted to briefly explain the multi-criteria decision-making process that takes into account the quantitative preferences to select the final conciliated treatment. Section 7 presents a case-study based on two real clinical guidelines for Diabetes Mellitus and Arterial Hypertension. The next section discusses the main limitations of our approach and the last section concludes.

2. Related Work

There exist diverse approaches to comorbidity that feature different models for the identification of merge criteria, representational schemes of CIGs or automation level of the reasoning procedures.

The automated conciliation in [13] proposes a divide-and-conquer strategy that distinguishes between clinical actions to measure the seriousness of the disease and treatment recommendations. Whilst this work is founded on the analysis of five clinical practical guidelines analyzed by a senior practitioner who designs a treatment model for each disease, other approaches rely upon experienced practitioners in comorbidity to identify the merging criteria and develop a representation ontology that captures the criteria to achieve the merging of multiple CPGs [14].

Handling interactions between comorbid diseases is of paramount importance in the development of automated tools for clinical decision support. In general, existing approaches put the emphasis in the design of automated reasoning processes to identify the most relevant interactions managed by clinicians [10]; namely, (1) redundant recommendations in more than one CPG; (2) undesirable physiological interactions between different drugs or between a drug and a disease; (3) scheduling constraints due to resource or temporal conflicts; or (4) contradictory recommendations. Techniques to address clinical interactions vary from their level of automation. Thus, one can find model-based automatic combination of multiple treatments [13], frameworks for eliminating redundancy and identifying adverse interactions [15], semi-automatic detection of interactions among redundant drug recommendations that require some attention from experts [16] or mixed-initiative proposals that provide physicians with management and reasoning tools when facing multiple CPGs [17]. A recent work suggests the introduction of different interactions strengths to measure the relevance of an interaction according to the features of the related elements [18].

The community of Artificial Intelligence in Medicine has already addressed the problem of planning comorbid patient treatments from different perspectives, which can be categorized depending on the outcome of the tool. This outcome can be either a *single multiple-disease CIG* that is dynamically created and then interpreted at execution time ([14], [17]), or it can be a personalized *multi-disease treatment plan* that is later delivered and enacted as a clinical pathway. Among the works that return a single multiple-disease CIG, the work in [13] applies an automated merging process of several treatment models (each obtained from the respective disease CPG) based on the pairwise combination of clinical actions and treatment tables through merging operators. The merging process is complemented by a rule-based execution engine that solves drug-interactions among actions. Other works, however, explicitly deal with multi-disease treatment plans, like the care plan oriented approach based on a semi-automatic knowledge acquisition process [19]. This approach is guided by the initial definition of a high-level and abstract plan previously identified by clinicians as the common steps to guide the development of the comorbid care plan. The process proposed in [19] is based on PROforma [20] and it is supported by knowledge engineering tools that help define a common guideline to concurrently treat two diseases by reusing elements in a plan library. The merging of guidelines is not fully automated and requires human intervention at every step in the life cycle. The work in [21] uses an ontology-based approach to merge Protégé-OWL *clinical pathways* of comorbidities, which are separately modeled and semi-automatically merged using the ontology merging techniques provided by Protégé. Merging comorbidity pathways ontologies involves solving duplications and managing drug interaction or adverse events since pathways (or the CG from which they are originated) may not explicitly state interactions. Other approaches focus on drug-drug and disease-drug interactions that arise in contradictory single-disease treatments, or suggest the utilization of specific programming models to analyze and detect interactions [22], which results in a combination of individual therapies derived from two involved CPGs. This latter work extends the proposal in [23] by using iterative actions for the administration of pharmacological treatments as well as numerical variables for the dosages in the administration of a medication.

The strength of agent-based representations has been exploited in projects like GLINDA [24] for solving interactions and consolidating treatment recommendations as well as in the K4Care platform, where agents personify the domain actors necessary for the execution of personalized home care treatments [25]. However, agent-based reasoning can be exploited not only for conciliation of medication in a reactive way, but also for planning tailored single-disease treatments following a deliberative approach. In previous works [26, 12], deliberative temporal HTN planning was used to represent formal CPGs[1], generate personalized care pathways and adaptively execute them [30].

The field of Multi-Agent Planning (MAP) is experimenting great advances over the last years by its capacity

[1]This representation can be done either directly, in the proper HTN formalism, by using a specific knowledge engineering tool [27, 28], or

and flexibility of modeling applications that require distribution of knowledge sources and reasoning. Particularly, the existence of disease specialist agents makes MAP a very suitable technique to address the comorbidity problem. MAP emphasizes the combination of two principal activities, planning and coordination. Some approaches apply a *pre-planning coordination*, distributing the task goals across agents before planning. This approach fits well tasks whose solution plan is made up as a composition of the individual plans of the agents [31]. MAP frameworks that apply *post-planning coordination* are particularly aimed to merge individual agent's plans into a single joint plan while solving interactions among the plans, as for example the partial global planning framework [32], which allows agents to communicate their local plans to the rest of agents and then merge this information into their own partial global plan in order to improve it. Finally, approaches that continuously interleave planning and coordination are more appropriate for tasks that require cooperative goals and where agents are not able to attain the goals by themselves [33].

In this work, we adopt a pre-planning coordination view differently to the classical MAP approaches. Rather than giving each single-disease specialist agent the goal of finding a plan (treatment) for its respective disease and then merging all the individual plans, agent coordination in our approach involves agents exchanging and combining the preference-filtered decision procedures of their respective CIGs. Our MAP approach allows each specialist agent to put forward its personal view in the conciliation of CIGs, thus obtaining as many patient-tailored agent-personalized conciliated CIGs as number of specialists involved in the comorbidity setting. Consequently, an equipoise situation can arise having as many treatment plans as conciliated CIGs. In this case, the patient treatment preferences are subsequently applied to select the most adequate multi-disease personalized and conciliated treatment.

As commented above, there are hardly approaches that deal with patient preferences in comorbid settings, with the notably exception of the work published in [34]. This work extends a first-order logic based framework to add support for patient preferences when mitigating the concurrently application of multiple CPGs to a comorbid patient. As we will see in the rest of the paper, our proposal brings some novelties with respect to the approach in [34] as a distributed representation of CPGs and a reasoning model that relies on a multi-agent coordination method. This provides a great flexibility to include or eliminate diseases as well as integrating new patient preferences without need to modify the reasoning routines.

3. Encoding clinical knowledge

The reasoning process of the agents within the proposed MAP architecture extends the temporal HTN planning process described in [35], which is used for the generation and execution of a personalized single disease care pathway in [26, 12, 36]. Based on the automatic translation of CIGs into HTN domain proposed in [12], this section shows an outline of the underlying knowledge representation and its adequacy to represent CIGs with patient preferences.

Most CIG languages are based on task-network models, characterized by a hierarchical decomposition of the care processes into networks of component tasks that unfold over time [6]. One of the most referenced CIG languages is Asbru, a task-network CIG language that emphasizes the representation and management of temporal aspects, specification and execution of action, hierarchical plans and parallel tasks [37]. While CIG languages like Asbru and HTN planning languages have shown to be very expressive in terms of their representation of temporal constraints [35, 38], the associated inference engines of CIG languages do not provide support for the automated generation of

trough an automated translation process from standard CIG languages to this formalism [12]. Moreover, this technology has been used to develop a pediatric oncology focused CDSS as a commercial application [29]

temporal plans. The work in [12] presents an automatic translation of CIGs specified with Asbru into the Hierarchical Planning Description Language (HPDL) that is used in this paper to formalize CPGs as HTN planning domains. The translation relies on an elicitation process for clinical guidelines described by experts in natural language which are then modeled with the user-friendly tool DELT/A (the Document Exploration and Linking Tool [39]) to produce a computer-interpretable XML-based model of the guideline. The proposal in [12] not only shows that HPDL is as expressive as other CIG languages, but also that HPDL can be used as an instrumental language that operationalizes clinical knowledge elicited with general tools like DELT/A. In short, HPDL enables knowledge to be interpretable and actionable by other software components such as planning systems, thus opening up the way to reason about knowledge and synthesize clinical plans.

An HTN planning domain described in HPDL is a compositional hierarchy of tasks that can represent either clinical procedures or goals to be achieved. Every task has associated (one o more) decomposition methods that describe alternative ways to accomplish it. A method describes the conditions under which a task can be decomposed (by means of a precondition) as well as how the task can be accomplished (by means of a partially ordered sequence of subtasks). Leaves of the hierarchy are called primitive actions and represent executable tasks that transform the world when executed. An HTN problem consists of an initial state (describing the health conditions of a patient, among other things) and a (partially) ordered sequence of high-level tasks to be accomplished (describing the intent to generate a plan for the diagnosis and treatment of a patient). An HTN planner receives as input a planning domain and a problem in order to generate a plan, i.e., a (partially) ordered sequence of primitive tasks. Previous work [12, 26] shows that HPDL allows for a full representation of a single-disease CIG as an HTN domain (mainly due to the capability of representing common workflow and temporal patterns required in CIGs [40]), and that this representation is directly interpretable by an HTN planner like [35] or [28] in order to obtain a full treatment for a single patient.

The simplest structure that can be found in the activities (or tasks) of a method, either clinical or administrative, is in sequence, within parentheses (see Figure 1) as a pattern with variables like `?patient` which generalizes the concrete names of patients or resources that take part in the activity.

```
((electrocardiogram ?patient ?room)(urine-analysis ?patient))
```

Figure 1: Representation of a sequence of activities

Activities may also be represented in parallel, within brackets, meaning that they can be carried out in any order, even simultaneously (see Figure 2).

```
[(eat-fruit ?patient)(isometric-exercise ?patient)]
```

Figure 2: Representation of parallel/independent activities

Each activity is described with several fields which encode the clinical knowledge or restrictions associated to the activity like its duration or the conditions that must hold for the activity to be enacted (see Figure 3).

`(:durative-action electrocardiogram`	Name of the activity
`:parameters (?p - Patient ?r - Room)`	Objects that take part and their types
`:duration (= 15m)`	Expected duration of the activity
`:condition (overall (available ?r))`	Prerequsitite: the room must be available all the time
`:effect (has-ecg ?p))`	Record of the activity

Figure 3: Representation of the knowledge associated to a single activity

Goals are thought of as compound activities which can be decomposed in different ways, meaning that there are

different manners of achieving the same goal. In this case, every method for decomposing the goal may be qualified with the conditions that must hold, if any, for this decomposition to be eligible (Figure 4).

```
(:task TreatmentArterialHypert                  Name of the compound activity or goal
  :parameters (?p - Patient ?d - Drug)          Objects that take part and their types
  (:method low                                  First decomposition method
    :precondition (severityAH ?p low)           Condition to be hold for the first decomposition to be eligible
    :tasks (                                    Representation of the first decomposition
     [(NonPharmaTreatAH ?p)
     ((DecidePharmaTreat ?p ?d ClassAH1)
     (DoPharmaTreatAH ?p ?d))]))
  (:method moderate                             Second decomposition method
    :precondition (severityAH ?p mod)           Condition to be hold
    :tasks (                                    Decomposition
     [(NonPharmaTreatAH ?p)
     ((DecidePharmaTreat ?p ?d ClassAH2)
     (DoPharmaTreatAH ?p ?d))]))
```

Figure 4: Compound activity, or goal, with two decomposition methods

In HPDL, the planning problem comprises the patient profile as well as information of the patient preferences (which is detailed in Section 4). We use a predicate-logic formalism to represent this information with simple facts that follow the form

(<attribute> <object> <value>)

where <object> indicates the patient in question, <attribute> denotes the feature and <value> is the value of the feature for the patient. The <attribute> component is the predicate name and the other two components are the arguments of the predicate.

The patient model integrates a comprehensive list of attributes that is extracted from three different information sources: Electronic Medical Records (EMRs), Clinical Guidelines and clinicians, following different methodologies for each one. Most of the patient profile information is commonly stored in EMRs, including biometric and personal details of a patient such as the age, the genre or the ethnicity, as well as information about his/her medical condition; e.g. diseases suffered by the patient, variables and levels of measurement. Figure 5 shows various HPDL facts representing personal details, medical conditions and contextual information. Other attributes related to health conditions are mainly elicited from CPGs following a manual process by consulting the textual guidelines. For example, `(exercise Jane regularly)`, `(training-inhalation Jane YES)`, `(diet Jane healthy)`. All in all, the set of attributes is defined considering different sources of information, and it may vary depending on the used guidelines, but the attribute modelling is ultimately carried out through a knowledge engineering process prior to the system is put into production.

By using the HPDL decomposition methods, facts related to patient profile attributes and medical conditions are quickly made actionable as it is shown in the compound activities of Figure 4, wherein the method `low` is activated when the fact `(severityAH Jane low)` is found in problem description, in which case the corresponding tasks are included in the treatment. Likewise, we can assume we have the facts `(age Jane 65)`, `(has-disease Jane Asthma)` in Jane's profile and the fact `(training-inhalation Jane YES)`, as described in Figure 5. This medical condition is interpreted here as a boolean value associated to a statement or observation of the patient, which may be a supportive piece of information to the clinician. Moreover, patient medical conditions and contextual information are used in methods' preconditions as a way to personalize treatments. Particularly, the administration of a treatment for

```
(:init
  (sex Jane M)                                  Genre of the patient
  (age Jane 65)                                 Age of the patient
  (has-disease Jane Asthma)                     Jane suffers from Asthma
  (has-disease Jane AH)                         Jane suffers from Arterial Hypertension
  (severity-AH Jane low)                        The severity of Jane's AH is low
  (bad-cholesterol Jane 221)                    Jane's bad cholesterol level
  (group Jane 1)                                Jane received an oncology protocol of group 1
  ...
  (exercise Jane regularly)                     Jane regularly does exercise
  (diet Jane healthy)                           Jane's diet is healthy
  (training-inhalation Jane YES)                Jane has experience in inhalation techniques
  ...)
```

Figure 5: HPDL facts that represent the patient profile and preferences.

Jane based on inhaled corticosteroids turns out to be appropriate so the system will apply the method corresponding to the inhaled corticosteroids treatment of the Asthma CPG, represented as an HTN domain in Figure 6.

```
(:task TreatmentAsthma
  parameters (?p - Patient ?d - Drug)
  (:method with_training                        method applicable when
  :precondition (training-inhalation ?p YES)    Jane has inhalation expertise
  :tasks (DoPharma ?p CorticoSteroids))         Generate a treatment with CorticoSteroids
  (:method no_training                          this method is applicable when ...
  :precondition (training-inhalation ?p NO)     Jane hasn't inhalation experience
  :tasks (DoPharma ?p AntileukoTrienes)))       Generate a treatment with AntileukoTrienes
```

Figure 6: An example of how treatments are personalized to the medical or contextual conditions of the patient.

The representation of time and expressive temporal constraints between activities, either compound or not, is also allowed in HPDL as shown in [35]. Every activity and goal is associated to three variables named `?start`, `?end` and `?duration` which delimit their temporal extent.

```
(:task TreatmentAH                              A compound task (or activity) to generate a treatment for Arterial
                                                Hypertension
  :parameters (?p - Patientm ?d - Drug)
  (:method low
    :precondition (severityAH ?p low)
    :tasks (
     [(NonPharmaTreatAH ?p)
     ((DecidePharmaTreat ?p ?d ClassAH1)
      ((= ?start (start A2))(DoPharmaTreatAH ?p ?d)))]))   Use of temporal constraints; ?start is the start time of task
                                                           DoPharmaTreatAH
  ...
```

Figure 7: Representation of temporal constraints (in bold faces) between activities: Activity `DoPharmaTreatAH` must be scheduled to start at the same time than activity `A2` regardless where it is within the plan.

For example, Figure 7 is a variant of Figure 4 which shows how the activity `DoPharmaTreatAH` is forced to start exactly at the same time than another activity named `A2`, which might appear in another branch of the plan. The `?start` and `?end` variables can be used from one activity to another to define a large variety of temporal relations, e.g. the ones described in Allen's temporal logic. For example, a temporally constrained task or goal can be written as `(<temporal_constraint> <subgoal>)`, and the standard Allen's temporal relation `OVERLAPS(g1, g2)` is represented in HPDL as

```
[g1 ((and (> ?start (start g1)) (< ?start (end g1)) (>= ?end (end g1))) g2)]
```

where ?start and ?end are variables representing the start and end points of g2, and (start g1) and (end g1) represent the start and end points of g1. The temporal planning process is described later in Section 5.5.

The patient model of the planning problem also encodes temporal facts representing the evolution of the patient profile along time that can be used to describe relevant changes in the patient conditions. Figure 8 shows that Jane's glycated hemoglobine (AC1) has changed from 7.0 on March 30 to 8.2 on June 28, what represents a significant change of Jane's medical conditions that affects a specific disease (Diabetes in this case). We can also observe in Figure 8 that other historical clinical aspects like the current therapy of the Jane are also encoded as temporal HPDL facts in the initial state. Specifically, Jane was prescribed to take Metformin 500 oral mode twice a day during three months since March 30. These two pieces of information, facts that represent the patient evolution and her current therapy, are crucial to find a personalized treatment adapted to the evolving state of the patient along time. Hence, the planning process (as explained in Section 5.5) will detect the change in the temporal evolution of Jane and will synthesize a new treatment adapted to the new health conditions of Jane. Otherwise, in case that Jane's medical conditions do not change, the current plan (therapy) will be adopted. Section 5.5 describes in detail this temporal planning process and Section 7.3 shows a detailed example about these concepts.

(:init	
...	
(at "30/03/2018 09:00:00" (AC1 Jane 7.0))	Jane's glycated hemoglobin was 7.0 on March 30
(Treatment Metformin 500 oral solution twice_a_day)	Jane's is receiving a treatment of Metformin 500, mode oral, format solution, at a daily frequency of twice a day
(start_Treatment Metformin "30/03/2018 09:00:00")	Jane's treatment began on March 30
(duration_Treatment Metformin 2160hrs)	Jane's treatment duration is 3 months (2160 hours)
(FollowUp Diabetes)	Jane's FollowUp for Diabetes
(start_FollowUp Diabetes "28/06/2018 09:00:00")	Jane's FollowUp was scheduled for current date
(duration_FollowUp Diabetes 1hr)	Jane's FollowUp session is estimated to last 1 hour
(= (start-this-visit) "28/06/2018 09:00:00")	A numerical predicate storing current visit's date
(at "28/06/2018 09:00:00" (AC1 Jane 8.2))	Jane's glycated hemoglobin is 8.2 at the current date
...)	

Figure 8: HPDL temporal facts representing the evolution of the patient status along time.

Finally, we introduce here the notion of *Decision Procedure* which will be very relevant in the forthcoming sections. As in the case of CIGs, the full HTN planning domain is not applicable to a given patient. The application of the patient condition and preferences in the domain results in a set of *Decision Procedures*, which comprise the activities of the decomposition methods that are eligible for the patient. For example, considering the task of Figure 6, a Decision Procedure for treating Jane's asthma would comprise only the method applicable to her training inhalation expertise.

In this paper, a multi-agent planning framework is defined so that every agent has the same planning capability but has a private HPDL-based knowledge representation that encodes a clinical guideline for a different disease. In the case of comorbid patients, these multiple planning agents collaborate with each other, by exchanging pieces of their private knowledge representation, that is to say, *Decision Procedures*, to coordinate the whole treatment. The multi-agent planning framework acts by merging the planning domain of every single disease, prior or during the search for a coordinated plan, instead of just merging the individual plans of the agents, which would amount to combining local treatments.

4. Patient preferences

In this work, we identify two types of preferences: *qualitative* preferences used by clinicians to make decisions while designing a personalized treatment, and *quantitative* preferences to facilitate the patient participation in the selection of a final treatment from many possible choices.

Handling patient preferences involves three main tasks: elicitation, representation in the proposed formalism and utilization in the decision-making process. Our particular interest is to incorporate patient preferences in the HTN planning process and, specifically, use the preferences to guide reasoning during conciliation of CPGs within the MAP framework. The issue of preference elicitation is out of the scope of this work and we will assume a manual elicitation for extracting both qualitative and quantitative preferences.

4.1. Qualitative preferences

Qualitative preferences are personal factors, desires and preferences that determine the patient habits and lifestyle. These are preferences to support the clinician during the development of the CPG. Qualitative preferences are elicited by the physician through an ordinary conversation to inquire the patient about guideline-specific preferences and to extract targeted information.

Specifically, a qualitative preference is represented as a HPDL fact such that a HTN method that encodes a particular action compliant with the preference will be triggered during the reasoning process. Similarly to the proposal in [34], we might obtain the following derivation when assessing the preferences of a patient, `Jane`, whose current treatment for diabetes (oral administration, diet and exercise) is not enough to keep the blood sugar level controlled:

1. *Oral medication* vs. *Injectable medication*. The user will respond according to the treatment benefits, burdens and side effects of both choices. Let's assume `Jane` prefers to keep on oral medication, which leads to a preference of the type `(medication-mode Jane oral)`.

2. *Side effects* vs. *Side effects*. Then, the physician presents the following two choices: (1) drugs which do not cause low blood glucose but has side-effects such as bloating, upset stomach or diarrhea vs. (2) drugs that can cause low blood glucose but has less severe side effects such as occasional skin rash or irritability. Let's assume `Jane` rather avoid the risk of low blood glucose. This is translated into health-related preferences like `(medication-type Jane Biguanides)` and `(medication-type Jane Alpha-glucosidase)`.

3. *Once a day* vs. *Twice a day* vs. *Each meal*. The next inquiry is about her preferable dosage schedule: taking pills at every meal, twice a day with breakfast and evening meal or take once a day in the morning. Assuming that `Jane` opts for taking pills once a day in the morning, the preference `(medication-time Jane breakfast)` is generated.

The above preferences activate the first method of the task for diabetes oral medication administration shown in Figure 9 which recommends *Metformin extended release* for a patient with the same preferences as Jane.

Another example of preference is a patient that suffers from Arterial Hypertension and opts for a treatment that reduces hospitalization rate although it causes persistent dry cough and dizziness (e.g., ACE inhibitors), over a treatment with diuretics that is less effective but has also less troublesome side effects such as frequent urination.

Qualitative preferences are formally represented as statements that are used to select the decision procedures that agents will handle at planning time. It is important to remark that the HTN tasks are usually decomposed in

```
(:task PharmaTreat-diabetes-oral
  :parameters (?p - Patient)
  (:method m1
    :precondition
      (and (medication-type ?p Biguanides)
           (medication-time ?p breakfast))
    :tasks (
     (:inline () (assign ?NRep (ndays)))
     (= ?start (breakfast-time ?p))
     (DoPharma ?p Metformin-extended-release ?NRep 24h))
  (:method m3
    :precondition
      (and (medication-type ?p Sulfonylureas)
           (medication-time ?p twice-a-day))
    :tasks (
     (:inline () (assign ?NRep (ndays)))
     (= ?start (breakfast-time ?p))
     (DoPharma ?p Glipizide ?NRep 12h)))
```

```
(:method m2
  :precondition
    (and (medication-type ?p Biguanides)
         (medication-time ?p twice-a-day))
  :tasks (
   (:inline () (assign ?NRep (ndays)))
   (= ?start (breakfast-time ?p))
   (DoPharma ?p Metformin ?NRep 12h)))
(:method m4
  :precondition
      (and (medication-type ?p Alpha-glucosidase)
           (medication-time ?p each-meal))
  :tasks (
   (:inline () (assign ?NRep (ndays)))
   (= ?start (breakfast-time ?p))
   (DoPharma ?p Acarbose ?NRep 8h)))
```

Figure 9: Oral medication for diabetes according to patient preferences

several methods and that a method embodies the line of reasoning to adopt when a personal preference, or simply a health-related condition, is fulfilled. However, if contradictory interactions are found at planning time with the preferences specified by `Jane` for the Arterial Hypertension guideline, an alternative method specified in the `PharmaTreat-diabetes-oral` task shown in Figure 9 or in the tasks of the Arterial Hypertension CPG will be launched.

4.2. Quantitative preferences

Preference (plan feature)	**Description**	**Label**	**Norm. Priority**
times_per_day	Counts dosage schedules weighted by dosing frequency	moderate (3)	0.21
pills_per_day	Counts the number of pills that the patient takes in a day.	medium (2)	0.14
number_of_medical_visits	Number of medical visits for the patient, including primary care, medical laboratory, hospitals, etc.	very much (5)	0.36
medical_visits_duration	Number of hours the patient must spend in all of the scheduled medical visits along the treatment.	moderate (3)	0.21
treatment_duration	The total duration of the treatment measured in days which gives an idea of the length and pacing of treatment, including, if necessary, durations of hospital stays. In order to label the criterion, consider that the shorter the duration, the better for the patient.	indifferent (1)	0.07

Table 1: Example of patient quantitative preferences

Quantitative preferences include measurable terms that are used to determine the patient perception over a treatment[2]. Quantitative preferences help select one option from many possible treatment outcomes (plans) and support shared decision making. In our approach, quantitative preferences capture the priority of a patient for a plan or treatment. These preferences are elicited through a simple and comprehensive questionnaire that enables the patient to respond questions about a given treatment. Patients provide answers in a qualitative scale, which is more suitable to answering certain kinds of questions, and then responses are automatically converted to a numeric rating scale. Table 1 shows the five patient quantitative preferences (plan features) we deal with in this work (first column) along with their description (second column).

The third column of Table 1 shows an example of the qualitative score of each preference given by a patient to a particular plan. Each row maps a preference to a qualitative score (label) representing to which extent the value

[2]Examination of costs from the institution standpoint is also addressed with quantitative preferences. This will be detailed in Section 6

of such preference in the plan is bothersome for the patient. The qualitative scores filled out by the patient describe the importance in terms of nuisance {*nothing, indifferent, medium, moderate, much, very much*} of each preference in the treatment or plan. The qualitative scores are mapped into a numerical value in the scale [0..5] since there are six possible values (the numerical values are shown within parenthesis besides the qualitative label). Then, we apply a normalization by the sum of all values to quantify the relative importance of each plan feature to the patient. Hence, a high priority for a preference denotes that the patient would prefer a conciliated plan with lower values of such feature. Following the values of the fourth column of Table 1, the most bothersome aspect of the plan to the patient is the number of medical visits, which means the patient would rather a conciliated plan with few medical visits. Details of the process that calculates the values and weights of the preferences are provided in section 6).

5. A MAP Architecture for conciliating CIGs accounting for patient preferences

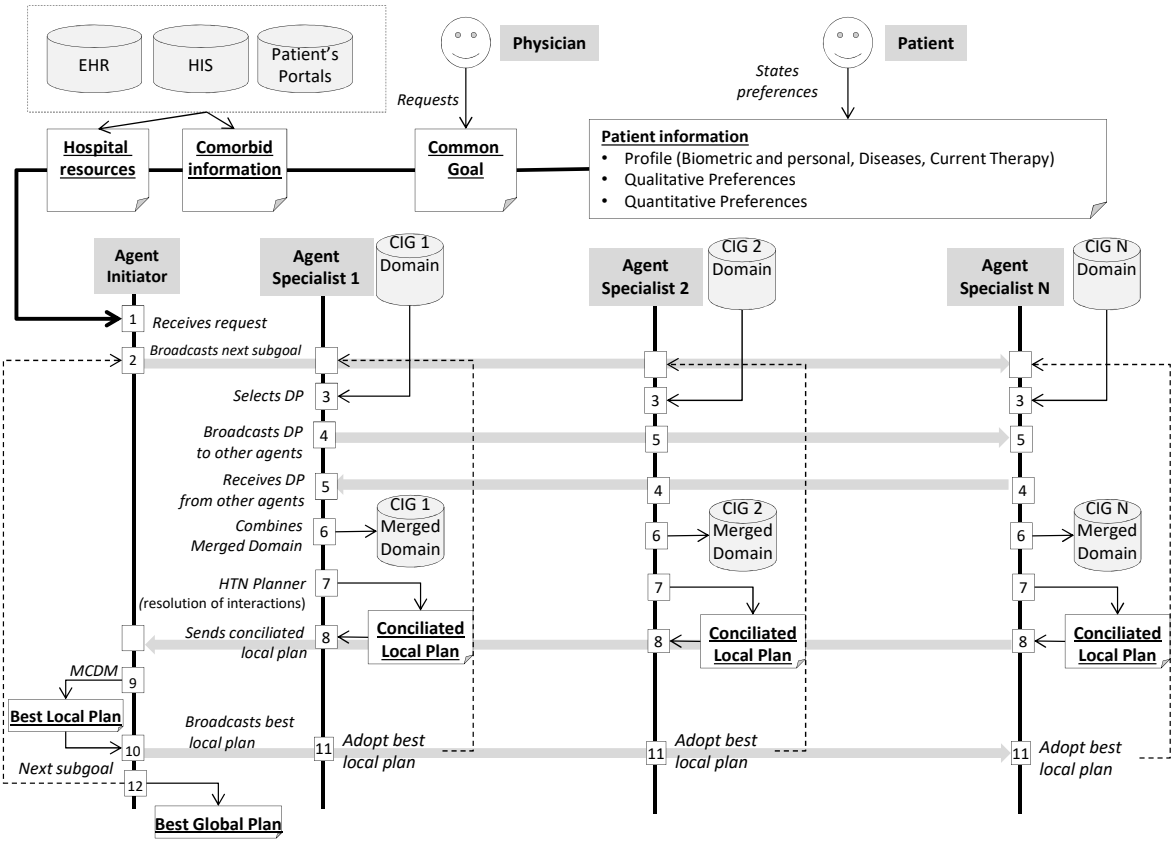

Figure 10: Multi-agent Architecture to conciliate *N* different diseases with an *Initiator* agent plus one *Specialist* agent for every disease

.

This section describes our proposed multi-agent planning system architecture (Figure 10) along with the technical details. The numbered labels in Figure 10 are used for the ease of the explanation. The *Initiator* agent receives

a request 1, which may come from a physician, to design the treatment for a comorbid patient. This request is named *Common goal* since, irrespective of the number of diseases a patient might present, this goal is common to all the agents and it is locally interpreted by each agent. The *Initiator* agent gathers all the information in the patient profile and launches the whole multi-agent process, integrating the information coming from the *Specialist* agents. There may be as many *Specialist* agents as required by the comorbidities of the patient, each disease being attended by a different *Specialist* agent that encodes the respective CIG expressed as an HTN domain, as sketched in the previous section. Most data needed to feed the multi-agent architecture might be obtained by dumping the data contained in Electronic Medical Records (EMR), or Hospital Information Systems (HIS) or Patient portals. Data are comprised by the following items: *Common goal*, *Patient information*, *Comorbid information* and *Hospital resources*. A detailed example about information items is provided in Section 7, Table 3. Note that in Figure 10 patient's profile also includes information about the *Current Therapy* followed by the patient at the beginning of the patient-clinician encounter (should it eventually have been prescribed in previous visits). It is a key piece of information to address dynamic aspects of the problem and to cope with the evolution of the patient, as described in Section 5.5 and shown in Section 7.3.

5.1. Common Goal management 2

The common goal takes the form of a set of (partially) ordered compound tasks (also called subgoals) that can be annotated with temporal constraints as the ones used in the decomposition methods (Figure 7). Subgoals are ordered as in decomposition methods; e.g., `(g1 g2)` stands for `g2` to be sequenced after `g1` and `[g1 g2]` means that both subgoals have to be scheduled in parallel. The *Initiator* agent receives the request to start a new planning episode and iterates over every single subgoal of the *Common goal*. In our case, the common goal of treating a patient is composed of two sequentially ordered subgoals: *diagnose* the patient and prescribe a *treatment*. The initiator sends the subgoals, in their corresponding order, to all the *Specialist* agents to be processed 2. The order of solving the subgoals is not mandatorily the same order to be followed at execution since additional temporal constraints may impose a parallel or any other time relation when executing the subgoals. In our particular case *diagnosis* and *treatment* will be executed in the same order as they are solved.

Subgoals are accomplished by following a planning process guided by the task decomposition schemes of an HTN domain. Each *Specialist* agent includes in its HTN domain a decomposition scheme (as shown in Figure 4) that specifies how to accomplish each subgoal of the *Common Goal*. This means that, each agent interprets every subgoal received from the *Initiator* differently and accomplishes it according to its own knowledge. For example, the accomplishment of the subgoal `((= ?start (current-visit-date)) (Treatment Jane))` could lead to different medication plans by different specialists. Moreover, tasks required to accomplish one same subgoal do not need to be synchronized nor aligned on time and can be executed independently of each other as long as they respect the temporal constraints specified in the domain (as explained in Figure 7, Section 3). An agent will not participate in the accomplishment of a particular subgoal if either its CIG does not directly contemplate the subgoal (in this case the agent would return an empty plan) or the conditions to accomplish the subgoal do not hold. For example, in the case of a patient who is already being treated for other diseases, an agent may decide not to accomplish the subgoal `(Treatment Jane)` of the Common Goal and maintain instead the patient's current therapy if the patient status has not significantly changed from the last visit (the underlying process is detailed later in Section 5.5).

The accomplishment of a subgoal starts the *conciliation of CIGs* as a local process composed of two steps, the *domain merging* and the *plan synthesis*. The domain merging step of a local agent is, in turn, composed of two sub-steps: filtering the Decision Procedures (DPs) to be sent to peer agents, and combining the DPs received from peer

agents with its local domain. Once agents have their own local merged CIG, they proceed to compute a plan with their embedded local planner. These processes are detailed below.

5.2. Filtering of decision procedures 3 4

Each agent carries out a local planning episode (with the local HTN domain and current patient state and preferences) aimed at *marking* the decomposition methods in every planning decision 3. Since patients' qualitative preferences are encoded in methods' preconditions, the marking process indeed filters the eligible decomposition methods according to the patient preferences, as well as other information included in the patient profile (see example of decision procedures in Figure 6). Once the planning episode ends with success, the marked methods are broadcasted to the peer agents 4. The sending agent is aware of the alternative methods that have been discarded, and this allows the process to backtrack and send the alternative decision procedures, in the case a concrete decision procedure would not be conciliable with the CIGs of other agents.

5.3. Combination of decision procedures into a merged domain 5 6

This process adds to the private local domain of each agent a renamed copy of the methods included in the decision procedures filtered by the other agents 5. This labelling step is required in order to later identify the actions of the conciliated plan that are applicable to each disease. Additionally, since the local planning process needs to know how to manage the decision procedures of the other agents, a *merged goal* is created too, for representing the simultaneous handling of the involved diseases 6. This merged goal takes the form of *parallel* compound activities including the local goals of each agent. For instance, suppose two different specialist agents *A* and *B* engaged in the resolution of the common subgoal `(Diagnosis ?p)`, where `?p` denotes the patient. Their local compound activities in the merged domain will look something like `(DiagnosisA ?p)` and `(DiagnosisB ?p)`.Therefore, the merged goal will take the form `[(DiagnosisA ?p) (DiagnosisB ?p)]`, that actually stands for *"apply concurrently two decision procedures, coming from different guidelines, to the same patient"*. This merged goal is used as the starting point for the HTN planning process, which is explained in the next section.

5.4. Plan synthesis: generation of a conciliated plan based on HTN Planning 7 8

Starting from the patient preferences and comorbidity information (represented in the initial state), along with the *merged domain and goal* provided by the merging process, an HTN planning process is carried out in order to generate a personalized conciliated plan 7 [3]. The planner follows a state-based forward search process, guided by the knowledge of the decomposition methods that were filtered according to patient preferences (and other patient contextual information). Higher level tasks are thereby decomposed in the order specified in the common goal. This means that once a goal has been accomplished by decomposing a high-level task, the planner keeps track of the resulting state after accomplishing that goal. Therefore, the search and reasoning process starts from, and accounts for, this updated state in order to achieve subsequent goals. Thus, the accomplishment of a goal is always aware of the resulting state after the application of a prior goal.

Moreover, besides *preference interactions*, the planning process also deals with the comorbidities of the patient and is able to manage *drug-drug, disease-drug, mutex, redundant intervention and multiple-support* interactions. In our approach, interactions are represented in HPDL in the public initial state (comorbidity information) as well as in the predicates and basic actions of the domain, which are assumed to be common to all the local guidelines.

[3] More details about how this planner is used for generating single-disease treatment plans can be found in [12].

```
(:durative-action admin-drug
:parameters (?p - Patient ?d - Drug
?ds - Dosage)
:condition (recommended ?p ?d))
:effects
(and (prescribed ?p ?d)
(when
(and (prescribed ?p ?other )
(drugs-interact ?d ?other))
(drug-drug-interaction ?p ?d1 ?d2)))))
(a)
```

```
Administration of a drug to a patient
with a dosage
The action is applicable if
the drug has been recommended in a
previous clinical decision
The drug is marked as prescribed
When
another prescribed drug
interacts with the current one ...
.. a drug-drug interaction is detected.

(b)
```

Figure 11: HPDL code for representing a drug administering action that detects if two drugs prescribed for the patient interact. The action makes use of HPDL standard concept of conditional effect, representing that some effects of an action are only applied when some condition holds in the planning state.

```
(:action isometric-exercise
:parameters (?p - Patient
?support - [discouraged recommended])
:condition ()
:effect
(and (recommend isometric-exercise ?support)
(when (and (recommend isometric-exercise ?other)
(/= ?support ?other))
(contradictory-support isometric-exercise)))
)
```

```
Doing isometric exercise may be
discouraged or recommended


Effects:
Record the recommendation
When another one exists
with a different support
Inform about this contradiction
```

Figure 12: HPDL code for representing an action that can arise *multiple-support* interaction.

Drug-drug **interactions** occur between two different drugs that must not be administered together. ***Disease-drug*** **interactions** occur when a drug must not be administered to a patient suffering a concrete disease. Such interactions can be represented with HPDL activities, as shown in Figure 11, and solved by following basic techniques used in planning as state space search.

Mutex **interactions** may occur when several actions, prescribed by different CIGs, require the same non-shareable resource and cannot be overlapped in time. HPDL can represent that a required resource must be available during the execution of an activity (see action `electrocardiogram` of Figure 3 wherein the availability of the non-shareable resource `?r` is checked in the action's preconditions). The local HTN planner resolves a mutex interaction by scheduling the interacting activity after the resource is released.

Redundant-intervention **interactions** occur when the same intervention is prescribed by different CIGs for the same patient, but should be applied just once to the patient. In order to represent this kind of interactions, HPDL allows a CIG designer to label some activities as *potentially redundant*. The local HTN planner detects an interaction when, once a potentially redundant activity is selected to be added to the plan, it finds another instance of the same activity in the plan. In this case the interaction is resolved by ignoring the action. The planner assumes that two instantiated actions that are to be added to a plan are redundant if they are instances of the same action pattern, are recommended by different guidelines and are labeled as *potentially redundant*.

Multiple-support **interactions** occur very often [1] when non-pharmacological redundant recommendations for comorbid patients (e.g. "do isometric-exercise") are recommended by different CIGs with contradictory support values (e.g., *recommended* or *discouraged*). This kind of interactions are represented as *redundant-intervention interactions*, but their effects allow to inform about contradictory recommendations (if any, see Figure 12). The local HTN planner detects an interaction when a non-pharmacological recommendation is redundant and its support value is contradictory with the same recommendation already included in the plan. The planner just solves the interaction by informing about the contradictory situation (e.g. "the recommendations about isometric exercise are contradictory").

***Preferences* interactions** occur when preferences cannot be met in the simultaneous application of two decision procedures coming from different CIGs. The local HTN planner detects a preference interaction when the application of an instantiated method (sent by another agent and merged into the local domain), *activated by a precondition representing a qualitative preference*, fails due to (part of) its activities cannot be conciliated with the local plan. This forces the local planner to backtrack and restart the conciliation process to request the sender agent an alternative decision procedure that meets the patient preferences. The case study in Section 7 shows a detailed example.

Once all the arising interactions are resolved, the output of the local HTN planner is a *patient-centered* (accounts for the patient preferences and clinical state) *multi-disease* (multiple CIGs are combined and served as a merged HTN domain) *conciliated* (interactions among CIGs are resolved) treatment plan for a comorbid patient. Since each *Specialist* agent gives priority to its own local decisions, and the patient preferences may be differently adapted by the agents, each *Specialist* agent might find a different conciliated plan. Then, the conciliated local plans are sent to the *Initiator Agent* 8 so that the most adequate plan from the patient standpoint is selected. A multi-criteria decision making process 9 (see the next section), wherein patient *quantitative preferences* play a crucial role, is performed to that end by the *Initiator* agent. Finally, the *Initiator* broadcasts the *best conciliated local* plan to every *Specialist* agent 10, and each one *adopts* the best conciliated local plan before proceeding with the next subgoal of the *Common goal*.

5.5. Plan adoption: addressing temporal and dynamic aspects of the problem 11

Plan adoption is a process by which an agent receives a plan as input and integrates it in its internal state. Plan adoption occurs in two situations:

- Within a planning episode when the *Initiator* sends the *Specialist* agents the best local conciliated plan that solves the diagnosis subgoal. Agents must adopt this conciliated plan and synthesize a treatment plan in accordance with the diagnosis plan.
- At the beginning of a planning episode when a *Specialist* agent opts for maintaining the current therapy of its disease because no significant change is observed in the health conditions of a patient during a follow-up visit. In this case the *Initiator* sends the rest of agents the plan they must adopt in their internal procedures.

The plan adoption of a *Specialist* agent involves a validation process that iterates on every action of the input plan using standard planning techniques to test if every action in the plan is consistent with the current plan under construction and the internal state of the planner. In the second case of plan adoption, *Specialist* agents do not build a new plan "from scratch" but they consider and integrate the plan of the agent that opted for maintaining its current therapy. Our MAP architecture is thereby able to cope with the dynamic aspects of the problem, accounting for changes in the patient status along time and considering as well a temporal planning horizon that determines to what extent the treatment plan should last.

Plan adoption as well as plan synthesis are addressed by a temporal planning process[4] that manages the temporal constraints described in the tasks of the HTN domain at different levels of abstraction (See Figure 13). Constraints are internally managed over a Simple Temporal Network (STN), a graph structure where nodes are time points bound to a domain of values and edges are the posted temporal constraints. During plan generation, a plan is deployed over the STN following a simple schema: (1) every task (either primitive or compound) is associated to a start and an end time point; (2) temporal points of compound tasks bound the time points of its subtasks; (3) temporal constraints are

[4]We refer the reader to [35] and [12] for a detailed description and experimentation about the specific single-agent temporal planning process.

`(:task TreatAndFollowup_AH_S1`	Task to schedule Treatment and FollowUp for Arterial Hypertension at the stage S1
`:parameters (?p - Patient )`	Applied to the patient ?p
`(:method adoptThePlan`	**Method for the case of no changing conditions**
`:precondition (not (changed-conditions ?p AH))`	If there are no significat changes in patient status
`:tasks ((AdoptPlan ?p AH)))`	Adopt the current treatment plan represented in the patient profile.
`(:method fromScratch`	**Method for the case of changing conditions**
`:precondition (changed-conditions ?p AH)`	If there are significant changes in Arterial Hypertension
`:tasks (`	Generate a new treatment
`(DecidePharmaTreat ?p ?d ClassAH1)`	Select an appropriate drug for this stage of AH
`((AND (> ?start (start-this-visit)` `(>= ?duration 3mths)` `(<= ?duration 6mths))`	*Temporal constraints for* `DoPharmaTreatAH`*: the treatment must start (variable* `?start`*) after the date of the current visit, and should last between 3 and 6 months*
`(DoPharmaTreatAH ?p ?d))`	Time points of all the subtasks of `DoPharmaTreatAH` are bound by the start and end points of this task
`((AND (>= ?start (+ (start-this-visit) 3mths))` `(<= ?start (+ (start-this-visit) 6mths)))`	*Constraints of this form provide flexibility for expressing the start time of a task execution. This constraint indicates that the next Follow-up session should neither start earlier than 3 months nor later than 6 months.*
`(FollowUp_AH ?p))`	In this method, every task prior to `FollowUp_AH` must end before its start point (planning horizon for the next visit) since they are totally ordered

Figure 13: An example of a task representing both, the conditions to carry out the adoption of a pre-existing treatment plan (method `adoptThePlan`), and how to schedule a pharmacological treatment (method `fromScratch`), with a duration between 3 and 6 months. The second method also schedules a follow-up session prescribed to occur between 3 and 6 months. The planning horizon is inferred by the temporal planning process following the temporal constraint satisfaction process explained.

encoded as absolute constraints with respect to the absolute start point of the STN; and (4) temporal constraints define the earliest and latest time of the start/end points associated to the tasks as described in Figure 13.

Whenever a compound or primitive task is added to the plan, all time points and constraints of the STN are posted, propagated and validated automatically (temporal consistency is based on the well-known AC3 algorithm [41]). This mechanism allows the planner to manage not only *explicit* temporal constraints derived from numerical constraints described in the domain, but also *implicit* temporal constraints derived from qualitative order constraints expressed in decomposition methods. The consistency of implicit temporal constraints involved in plan-subplan relations is also granted, since the time points of subtasks of any task t are bound and embraced by the time points of t; that is, *subtasks inherit the constraints of their higher-level task*. This makes it possible to represent and reason about temporal constraints derived from hierarchical decompositions.

In practice the temporal horizon of a plan is inferred by the planner by using this temporal HTN planning process. As shown in Figure 13, the planning horizon is indirectly derived by means of a process that accounts for the temporal constraints of the tasks. In this Figure, the high-level pharmacological treatment must end before the follow-up task starts, which in turn is scheduled to start by [3,6] months after the date of the current visit. This forces the planner to schedule all the actions prior to the follow-up session to occur before the follow-up task. Hence, the start time of the follow-up tasks plays the role of the planning horizon (see the example in Section 7.3 for more details).

Consistency checking of temporal constraints is a key operation in our approach since the temporal planning process is aimed to find legal executable treatment plans from a given input CIG. If an inconsistency is detected at planning time in the STN during the decomposition of a given task t, the search process backtracks and tries to find another valid decomposition for t, seeking alternative decomposition methods. The search process continues until the

space of alternative decompositions is exhausted. In the case that a temporally inconsistent guideline is provided as input, the temporal HTN model will also be inconsistent. No legal executable pathway can be found and, therefore, the planner will return a FAIL.

When the *Common goal* contains pending subgoals 12, the process iterates back to 2. Otherwise, the best aggregated plan is finally obtained and sent back to the physician and patient to support their final decision.

6. Selection of a conciliated plan accounting for quantitative preferences.

The conciliated plans or treatments proposed by each *Specialist* agent adhere to merged clinical guidelines and fulfill patient qualitative preferences as far as possible but may differ in characteristics like medication administration complexity, treatment duration or number of medical visits. The *Initiator agent* upon reception of the proposed conciliated plans sent by specialists, carries out a Multi-Criteria Decision Making (MCDM) process [42] that helps select the most suitable conciliated plan according to patient (resp. institution's) *quantitative preferences*[5].

The MCDM process scores every plan j received by the *Specialist* agent through a plan evaluation function, $p(j)$, that takes into account the patient quantitative preferences (see Table 1). This process requires calculating the value of the five preferences (plan features) in plan j and the weight or priority of each preference to the patient.

Values of the plan features. The values of the features in plan j are calculated as follows:

1. `times_per_day` quantifies the complexity of the medication regimen by summing the number of different dosage schedules weighted for dosing frequency [10]. For instance, if the treatment in j consists of m drugs taken twice per day and n drugs taken once per day, this results in a complexity score of 3 (2+1). A patient will most likely opt for a treatment with few drugs and few different dosage schedules.

2. `pills_per_day` complements `times_per_day` to account for the complexity of the treatment administration. In the above example, the value of this feature is $m + n$.

3. `number_of_medical_visits` is calculated by counting the actions of j that represent a medical visit.

4. `medical_visits_duration` is calculated by summing up the value of the parameter `:duration` of all the medical visit actions

5. `treatment_duration` is the plan duration

Since features are measured in different units, the values of the five features in plan j are normalized with respect to the values of the features in the rest of plans. Given a set of plans P, $1 \leq j \leq P$, the normalized value of feature i in plan j, n_{ij}, is calculated by considering the value of the same feature i in every plan of P: $n_{ij} = v_{ij} / \sum_{k=1}^{|P|} v_{ik}$.

Weights of the patient preferences. As explained in Section 4, the user qualitatively scores the five plan features in Table 1. Scores are then translated into a numerical rating scale. A numerical value of 1 is assigned to the lowest ranked criterion (the least annoying) and for each qualitative score up in the ranking, we increase the numerical value by one. For each feature $1 \leq i \leq 5$, the numerical value assigned to i (s_i) is divided by the sum of the numerical values assigned by the patient to the five plan features, thus yielding a normalized priority value for each i: $pr_i = s_i / \sum_{k=1}^{5} s_k$.

[5]The MCDM process is equivalent for handling institution's preferences. Details on this type of preferences are given later

Finally, the formula that returns how much the patient *dislikes* plan j is: $p(j) = \sum_{i=1}^{5} pr_i * n_{i,j}$.

The MCDM process can be easily extended to account for preferences of other actors involved in the care process. Medical institution preferences, for instance, are especially valuable by clinicians so as to promote medical decision making while considering treatment costs and optimization of resources. In this regard, quantitative preferences that reflect institution preferences can also be included with a new evaluation function, $m(j)$, that considers the costs incurred along the treatment. Table 2 shows the five quantitative institution preferences included in the evaluation of the treatment cost. The evaluation process is equivalent to the one described for patient preferences and the value of $m(j)$ is calculated with the same formula that computes $p(j)$.

In order to aggregate both functions, the final score for a treatment is computed as a weighted average $S(j) = w_p * p(j) + w_m * m(j)$, where the weights ponder the relative importance of each category of preferences. A deeper analysis of the patient and institution preferences and the utility functions might lead to discover nonlinear relationships, thus allowing for a wider range of possible dependencies.

In any case, once the plans are scored, the *Initiator* agent selects the plan with the lowest score $S(j)$. This automated selection process could be replaced with a mixed-initiative process where both clinician and patient interact. Subsequently, the plan is sent to and adopted by every *Specialist* agent in order to proceed with the rest of activities of the common goal.

Criteria	Description
drug_cost	The cost of the drug administration in the patient treatment.
laboratory_test_cost	The cost of blood tests, urine tests, body tissues tests,...
RX_study_cost	The cost of regular X-rays, MRI or CT scans.
medical_visit_cost	The cost of the medical visits to the specialist.
hospitalization_cost	The cost per impatient per day.

Table 2: Example of medical institution's quantitative preferences. The qualitative scores are filled out by a clinician and describe how relevant is a criterion with respect to the cost of a treatment {*nothing, indifferent, medium, moderate, much, very much* }. They are later mapped into a numeric value (within parentheses) that is translated into a normalized priority value.

7. Experimental validation

In this section, we present a case study to analyze the entire life cycle of the architecture while illustrating a comprehensive view of the concepts and procedures exposed in the prior sections. Section 7.1 shows the generation of the plans at the Diagnosis stage and section 7.2 explains the solving process of the Treatment stage. Following the case study, Section 7.3 presents four validation scenarios that exhibit the health conditions of a patient at different times as well as the progressive adaptation of the conciliated treatment plan for the patient. Finally, section 7.4 describes the insight of the clinicians about the behaviour of the tool.

The case study is based on an architecture configured with two *Specialist* agents (labeled *A* for AH and *D* for DMT2) and one *Initiator* agent. The information required is summarized in Table 3. The HPDL representation shown in the table is a simplification devoted to illustrate only the main issues described so far. In the remainder of this section, we will illustrate that during the accomplishment of the Diagnosis subgoal, only mutex interactions and redundant intervention interactions arise. Then, during the accomplishment of the Treatment subgoal, drug-drug, drug-disease and multiple-support interactions arise. Qualitative preferences interactions arise during the accomplish-

Patient profile	
This information is represented as HPDL conditions, as shown in Section 3 Figure 5.	The case is represented by a 65-year-old female named *Jane*, who has a Systolic Blood Pressure (SBP) of *145*, a comorbid patient with Diabetes Mellitus Type II (DMT2) and Arterial Hypertension (AH).
Quantitative preferences	
Patient preferences are shown in Table 1, Section 6. Medical institution's preferences are shown in Table 2, Section 6.	In a initial step, Jane fills out a form to capture her quantitative preferences. The clinician fills a similar form to capture medical institution's preferences.
Qualitative preferences	
`(medication-time Jane breakfast)`	Jane prefers to take the medication at breakfast time. The value of this qualitative preference is provided by the clinician after asking the patient for her preference about taking medications.
Comorbitiy information	

`(drug-is-for-disease GLI DMT2) (drug-is-for-disease CHL AH)`
`(drug-is-for-disease GLY DMT2) (drug-is-for-disease TRA-VHC AH)`
`(drug-is-for-disease ACA DMT2) (drug-is-for-disease TEL-AML AH)`
	Recommended drugs for each disease. GLIpizide, GLYburide and ACArbose are drugs for DMT2 and CHLorthalidone, TELmisartan + AMLodipine and TRAndolapril + VerapamilHydroChloride are drugs for AH. These drugs have been represented after consulting the corresponding CPGs. These facts are needed to select which drugs can be prescribed or to decide alternative drugs to resolve drug interactions.

`(drug-admin-time GLI breakfast)  (drug-admin-time GLY breakfast)`
`(drug-admin-time CHL breakfast)  (drug-admin-time TEL-AML breakfast)`
`(drug-admin-time TRA-VHC breakfast)  (drug-admin-time ACA each-meal)`
	Drugs' administration time is also provided, in order for the agents to both adapt pharmacological treatments to patient qualitative preference, and detect and resolve *qualitative preferences interactions*. This information is known by clinicians either by consulting the administration mode of drugs recommended by a CPG, or by their own clinical practice experience

`(drug-interacts-with-drug TRA-VHC GLI)`
`(drug-interacts-with-drug TRA-VHC GLY)`
`(drug-interacts-with-disease CHL DMT2)`
	Drug-drug or drug-disease interactions are represented with these facts, meaning that Trandolapril plus VerapamilHydroChloride interacts with both GLIpizide and Glyburide, while CHLorthalidone interacts with the disease DMT2. This information will be used to resolve drug-drug and disease-drug interactions as will be shown later in this section.

Resources	
`(available room1)`	This fact represents a non-shareable room for some clinical test. In order to evaluate quantitative institutional preferences, information about costs of actions related to diagnosis procedures have to be provided. We have assumed integer monetary costs for the following actions related to institutional criteria: ***laboratory_test_cost*** (i.e., *blood-test cost*: 4, *urine-analysis cost*: 3, *fasting-plasma-glucose cost*: 4.); ***radiographic_study_cost*** (i.e., *electrocardiogram cost*: 5); and ***medical_visits_cost*** (i.e., *measure-blood-pressure = measure-height-weight = measure-waist-circumf*: 0, *check-clinical-history cost*: 1, *check-organ-damage cost*: 1, *ocular-tests cost*: 1).
Common Goal	
`((Diagnosis Jane) (Treatment Jane))`	The conciliated plan should be built in two sequential steps: the first subgoal represents the step at which activities required for the patient diagnosis need to be included in the diagnosis plan and the second one represents that the conciliated plan should also embody the pharmacological and non-pharmacological treatment activities.
CIGs represented as HTN domains	
The CIGs represented for this case are built on two real CPGs for Diabetes Mellitus Type II (DMT2) [43] and Arterial Hypertension (AH) [44].	CIGs' representation has been carried out in collaboration with a team of three clinicians, two internists and one pediatric oncologist with several years of expertise, at the Hospital San Agustín of Linares (Spain), who also participated in the validation of the proof of concept (described at the end of this section).

Table 3: Required information to carry out the case study. Temporal facts about patient conditions and current therapy are addressed in examples of Section 7.3.

```
(:task DiagnosisDMT2 :parameters (?p - Patient)
 (:method diagDMT2   :precondition ()
  :tasks((blood-test ?p)
         (urine-analysis ?p)
         (fasting-plasma-glucose ?p)
         (measure-height-weight ?p)
         (measure-waist-circumf ?p)
         (measure-blood-pressure ?p)
         (ocular-tests ?p))))
```

```
(:task DiagnosisAH :parameters (?p - Patient)
 (:method diagAH   :precondition ()
  :tasks( (measure-blood-pressure ?p)
          (blood-test ?p)
          (urine-analysis ?p)
          (electrocardiogram ?p)
          (check-clinical-history ?p)
          (measure-height-weight ?p)
          (measure-waist-circumf ?p)
          (check-organ-damage ?p) ))
```

Figure 14: Decision Procedures filtered from local CIGs for the subgoal (Diagnosis Jane) and for the patient *Jane*: DMT2 (left) and AH (right). A suffix labeling the disease has been added to the names of compound tasks and decomposition methods.

ment of the Treatment subgoal too, and we will show that the cause of this interaction is that a qualitative preference cannot be fulfilled by one of the agents, which asks for alternative decision procedures to the other agent.

7.1. First subgoal: Diagnosis

Each Specialist agent selects the set of DPs from its local CIG that are suitable for this subgoal and Jane's qualitative preferences (see Section 5.2). In this case, the Decision Procedures of agents *A* and *D* are both composed of a single method of their respective CIGs (e.g., methods *diagDMT2* and *diagAH*, respectively, in Figure 14). Each *Specialist* agent sends its DPs to the other agent. Then, each *Specialist* agent combines the received DPs with its local CIG, proposing a merged domain for Diagnosis and a merged goal that represents the concurrent application of its own Diagnosis method with the received DPs (as explained in Section 5.3).

Note that many actions in Figure 14 are redundant, and the actions *electrocardiogram* and *ocular-tests* (both in different CIGs) make use of the non-sharable resource `room1` (as explained in Section 3, Figure 3). With this input information, each agent carries out the plan synthesis, coming up with a conciliated plan (see Figure 15). The conciliated plans of agents *A* and *D* are different because each agent gives priority to its local CIG when adding actions to the plan. That is to say, agent *D* first adds to the plan the actions of its own local method *diagnosisDMT2* (Figure 14), and then adds the actions of the DPs coming from agent *A*. While adding these actions, agent *D* tries to resolve *redundant-intervention* and *mutex* interactions (Section 5.4) with the actions recommended by *A* , but not all can be resolved (except the mutex) due to the ordering constraints imposed by the DP of agent *A* (thus, it simply appends these actions to its conciliated plan). On the contrary, agent *A* is able to synthesize a conciliated plan since it manages the subgoals of its local merged goal in a different order.

		Global Plan agent D			**Global Plan agent A**		
Criteria (i)	pr_i	v_{iD}	n_{iD}	$pr_i * n_{iD}$	v_{iA}	n_{iA}	$pr_i * n_{iA}$
laboratory_test_cost	0.33	18	0.62	0.2	11	0.38	0.13
radiographic_study_cost	0.33	5	0.5	0.17	5	0.5	0.17
medical_visit_cost	0.33	3	0.5	0.17	3	0.5	0.17
		$\sum_{i=1}^{3} pr_i * n_{iD}$		0.54	$\sum_{i=1}^{3} pr_i * n_{iA}$		0.47

Table 4: Medical institution criteria for Diagnosis stage

Once the conciliated plans are synthesized, agents send them to the *Initiator*, and the MCDM process is applied in order to evaluate and select the best plan (see Section 6) according to *patient and institutions quantitative preferences*

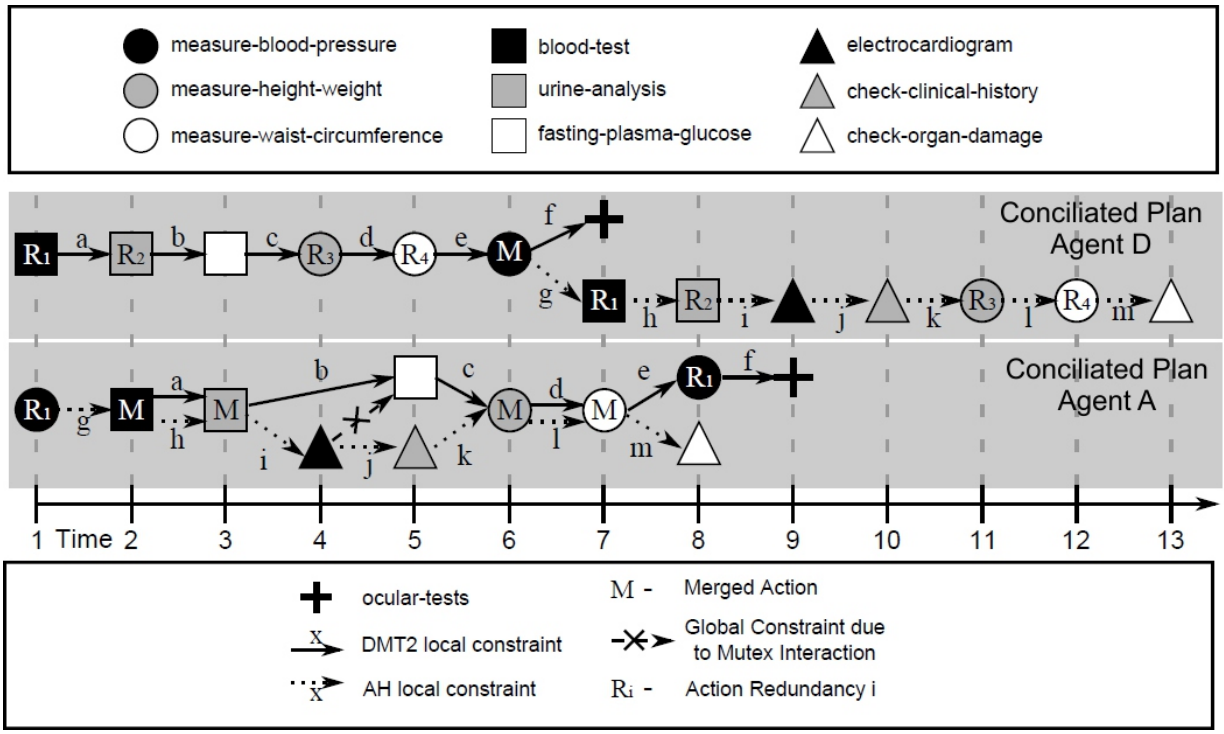

Figure 15: Conciliated Plans for the Diagnosis stage, synthesized by the local HTN planner of each agent. The conciliated plans contain information about the merged and redundant actions.

and (see Tables 1 and 2 in Section 6). In the Diagnosis stage of this case study, the only criterion that has an impact on the patient preferences is *number_of_medical_visits*. Because of that, the normalized values of the two conciliated plans are calculated by using only the *number_of_medical_visits.* The conciliated plan of agent D is composed of 14 medical visits (Figure 15, top), and the one of agent A is composed of 11 medical visits (Figure 15, down). If we normalize these values we get the values $p(Conciliated\_Plan\_D) = 0.56$ and $p(Conciliated\_Plan\_A) = 0.44$. Medical institution's criteria and values are shown in Table 4, assuming for this case study that only three criteria are considered and all with the same priority of 0.33. Finally, as explained in Section 6 the final value of the global plans taking into account quantitative preferences of both patient and medical institution is calculated as a weighted average: $s(Global\_Plan\_D) = 0.7 * 0.56 + 0.3 * 0.54 = 0.55$, and $s(Global\_Plan\_A) = 0.7 * 0.44 + 0.3 * 0.47 = 0.45$. In this case we set the weight of the patient preferences to 0.7 and the medical institution's preferences to 0.3, thus giving more importance to the patient opinion.

Therefore, the conciliated plan obtained by agent A is selected because of its lower score, and each agent adopts this plan before proceeding with the next subgoal.

*7.2. **Second goal:*** **Treatment**

The next (and last) subgoal of the common goal is `(Treatment Jane)`. The *Initiator* selects this subgoal and sends it to the *Specialist* agents. Agents filter its set of Decision Procedures and keep the methods for administering the pharmacological treatment at breakfast time, which is Jane's personal preference. Agents send the filtered DPs to each other and try to conciliate them with their own through their local HTN planner.

Table 5 shows that agent D is able to find a conciliated plan in this step, but agent A is not. On the one hand,

Specialist Agent	Non Pharmacological Treatment	DMT2 Pharma Treatment	AH Pharma Treatment
D	intake-alcohol(D)(Cons) eat-fruit(R)(Cons) physical-exercise(R)(Cons) isometric-exercise(R)(Contra) isometric-exercise(D)(Contra) glucose-monitoring(R)(NMS) smoke(D)(NMS) intake-sodium(D)(NMS) weight-reduction(R)(NMS)	GLI at breakfast (1 per day) during 10 days	TEL+AML at breakfast (1 per day) during 10 days
A	**First attempt to conciliate (preference interaction detected)**		
	no plan found		
A	**Conciliated plan after asking for alternative DPs**		
	intake-alcohol(D)(Cons) eat-fruit(R)(Cons) physical-exercise(R)(Cons) isometric-exercise(R)(Contra) isometric-exercise(D)(Contra) glucose-monitoring(R)(NMS) smoke(D)(NMS) intake-sodium(D)(NMS) weight-reduction(R)(NMS)	ACA once per meal (3 per day) during 10 days	TRA+VHC at breakfast (1 per day) during 10 days

Table 5: Conciliated Plans for the subgoal `(Treatment Jane)`. Agent *A* needs to backtrack and ask for alternative methods to conciliate its plan. The format for recommendations is *recommendation-name(local-support)(aggregated-support)*. The aggregated support value is set in the *aggregated-support* metatag of the actions. Abbreviations: R = Recommended; D = Discouraged; Cons = Consistent; Contra = Contradictory; NMS = No Multiple Support.

regarding the non-pharmacological part of the treatment, agent *D* is able to detect and manage all of the *multiple-support* interactions (Section 5.4) within the non-pharmacological treatments, so the aggregated support is reflected in the conciliated plan (for instance, the action *isometric-exercise* is found in both plans with local values of support *recommended* and *discouraged*, so it is labeled with the *contradictory* aggregated support value).

Regarding the pharmacological treatment, it is worth analyzing the result from the each agent standpoint:

- **Agent *D*'s standpoint:** agent *D* first adds the actions and constraints from its own CIG to the conciliated plan. The drugs are selected from the alternative instantiations of predicate *(drug-is-for-disease ?drug DMT2)*. As a result of this process, GLIpizide is the drug selected for treating DMT2 since it is the first drug declared in the initial state. This is a selection criterion that has been illustrated for the sake of simplicity in this example. Indeed, HPDL has the capability of using more complex expressions in actions and methods preconditions in order to use other selection criteria, for example using heuristic/utility values to select the most appropriate drug for the patient. The plan built so far is a sequence of instantiated drug-administering actions, each one represented as shown in Figure 11, Section5.4.

  Afterwards, the DPs received from agent *A* are processed, and Agent *D* tries to determine a non-interacting drug for Arterial Hypertension, using those DPs. According to initial state drug interaction facts, the only option is to prescribe TEL-AML as a drug to treat Arterial Hypertension. This way, Agent D is capable of obtaining a conciliated plan that incorporates a pharmacological treatment for the disease AH based on TELmisartan+AMLodipine. because there is no interaction among this combination of drugs and GLIzipide.

- **Agent *A*'s standpoint:** agent *A* does not find a conciliated plan because of a *preference interaction*. It first adds the actions and constraints from its local CIG to its conciliated plan and, according to its priorities, the pharmacological treatment is, in a first attempt, based on TRA-VHC. Afterwards, the DPs coming from agent *D* are processed, but no pharmacological treatment can be found because all the drugs for DMT2 (namely Glipizide and Glyburide) that *can be prescribed for breakfast* interact with TRA-VHC. At this point, agent *A* asks agent *D* for another alternative set of DPs to deploy a pharmacological treatment for DMT2 conciliable

with its local CIG (for this to be possible, *D*'s local CIG should have been appropriately designed with such an alternative). In this case, agent *A* is able to conciliate the new set of DPs coming from *D* by using Acarbose (to be taken by the patient not at breakfast but at each meal) for treating DMT2.

Table 5 shows the conciliated plan, wherein patient *personal preferences* are met as far as possible. Note also that the conciliated plans of both agents are substantially different with respect to pharmacological treatment, but both valid from a clinical standpoint. For this reason the MCDM process is needed to resolve this equipoise situation. Finally, it is worth mentioning that in the non-pharmacological treatment shown in Table 5 there are two actions that recommend isometric-exercise with contradictory support. The planner is capable of analyzing and simplifying redundant recommendation when they are not contradictory (as explained in Section 5, but it has not enough knowledge (i.e. we have not encoded knowledge in the CIG) to make an autonomous decision about which action should be recommended when two different guidelines make a contradictory recommendation. However, the planner is capable of detecting such situation and at present, the only proposed solution is to inform the clinician about this contradiction.

Considering the application of the MCDM process to select among the proposed conciliated plans, the patient quantitative preferences are influenced by the *times_per_day*, *pills_per_day* and *treatment_duration* criteria. However, in this case institution's preferences are not used, since none of the institutional criteria is applicable to evaluate Jane's conciliated treatment plan. Table 6 shows the priorities obtained from the form filled by Jane giving importance to each criterion, along with the value of each criterion in both plans as well as their normalized values. The value of Jane's quantitative preferences for both conciliated plans are obtained as explained in Section 7. Consequently, agent *D*'s proposal is preferable over *A*'s proposal due to its lower score. Finally, as described in Section 5 (points 8 and 12), the system returns agent *D*'s proposal as the best plan and, since there are no more pending goals, the process terminates.

		Global Plan agent D			**Global Plan agent A**		
Criteria (i)	pr_i	v_{iD}	n_{iD}	$pr_i * n_{iD}$	v_{iA}	n_{iA}	$pr_i * n_{iA}$
times_per_day	0.5	1	0.2	0.1	4	0.8	0.32
pills_per_day	0.33	2	0.33	0.1	4	0.666	0.22
treatment_duration	0.17	10	0.5	0.09	10	0.5	0.09
		$\sum_{i=1}^{3} pr_i * n_{iD}$		0.29	$\sum_{i=1}^{3} pr_i * n_{iA}$		0.63

Table 6: Patient quantitative preferences to evaluate criteria for the Treatment stage. Priorities's values ($pr[i]$) are obtained after normalizing the qualitative values provided by the patient for each criterion. **Moderate** for *times_per_day*, **Medium** for *pills_per_day* and **Indifferent** for *treatment_duration*.

Note that this is a step-wise process in which the conciliated plan is incrementally generated and synchronized by the *Initiator* agent. The high-level steps in the common goal of this case study are just an example, and should not be considered as the only strategy to come up with a conciliated plan. Indeed, it is possible to differently customize the common goal used as guide towards conciliation accordingly to the patient case.

7.3. Validation scenarios

We implemented a proof of concept of the architecture and life cycle described in Section 5. Our aim was to validate the technical feasibility of the described techniques, highlighting that MAP is an enabling technology to

Scenario (*t*)	Description (*t*)	TEMPORAL FACTS ABOUT PATIENT PROFILE(*t*)	
		Health Conditions at *t*	**Current therapy followed by the patient prescribed at scenario *t-1***
Scenario1 (30/03/2018) First visit	A 46 years old woman with A1C = 7, SBP = 120, DBP = 80, CVDRisk = 0.1	(A1C Jane 7.0) (SBP Jane 120) (DBP Jane 80) (CVDR Jane 0.1)	None (New patient)
Scenario2 (26/04/2018) Eventual visit because of elevated blood pressure	Patient keeps Diabetes conditions. Curent therapy metformin 500 twice a day. Diabetes Revision at 28/06/2018. Detected Elevated BP (131/82). CVDRisk = 0.2	(A1C Jane 7.0) (SBP Jane 131) (DBP Jane 82) (CVDR Jane 0.2)	(Treatment Metformin 500 oral solution twice_a.day) (start_Treatment Metformin "30/03/2018 09:00:00") (duration_Treatment Metformin 2160hrs) (FollowUp Diabetes) (start_FollowUp Diabetes "28/06/2018 09:00:00") (duration_FollowUp Diabetes 1hr)
Scenario3 (28/06/2008) Scheduled FollowUp for Diabetes and for Hypertension	Hypertension current therapy is Lisinopril 10 once a day. Hypertension conditions remains the same since last visit (26/04/2018). Diabetes therapy is Metformin 500 twice a day but at this FollowUp, after 3months, Diabetes conditions change from A1C= 7 to A1C = 8.0.	(A1C Jane 8.0) (SBP Jane 131) (DBP Jane 82) (CVDR Jane 0.2)	<same Diabetes current therapy than Scenario2> (Treatment Lisinopril 10 oral tablet once_a_day) (start_Treatment Lisinopril "26/04/2018 09:00:00") (duration_Treatment Lisinopril 720hrs) (FollowUp HyperTension) (start_FollowUp HyperTension "28/06/2018 09:00:00") (duration_FollowUp HyperTenmsion 1hr)
Scenario4 (30/09/2018) Scheduled FollowUp for Diabetes and Hypertension	Hypertension current therapy is Lisinopril 10 once a day. Hypertension conditions have evolved to Hypertension Stage2. Diabetes therapy is Metformin 500 twice a day + Glyburide 1.25 three a day. At this FollowUp Diabetes conditions change from A1C = 8.0 to A1C= 7.0.	(A1C Jane 7.0) (SBP Jane 145) (DBP Jane 92) (CVDR Jane 0.2)	<same Hypertension current therapy than Scenario3, though revision date changes> (FollowUp HyperTension) (start_FollowUp HyperTension "30/09/2018 09:00:00") (duration_FollowUp HyperTenmsion 1hr) (Treatment Metformin 500 oral solution twice_a_day) (start_Treatment Metformin "28/06/2018 09:00:00") (duration_Treatment Metformin 2160.hrs) (Treatment_second Glyburide 1.25 oral tablet three_a_day) (start_Treatment_second Glyburide "28/06/2018 09:00:00") (duration_Treatment_second 2160hrs)

Figure 16: A summary of the four scenarios set to validate the architecture. Column Scenario shows several visits of the patient at different times (from the earliest to the latest date). Column Description shows a commonly written text by a clinician to describe the current status of the patient (AC1 stands for glycated hemoglobin value, SBP for Systolic Blood Pressure, DBP for Diastolic Blood Pressure, and CVDRisk for Cardiovascular Disease risk). Column Health Conditions shows the HPDL temporally annotated facts that represents the health conditions when the patient comes to consultation. The last column contains the temporal facts that represent the current therapy the patient is taking when (s)he comes to the visit.

conciliate CPGs in comorbidity settings accounting for patients preferences. Additionally, our experimentation is designed to show that our architecture is capable of facing with dynamic situations in multiple scenarios, coping with the evolution of a patient whose health conditions change during the application of a treatment which, in turn, has to be modified and adapted to the new patient health conditions. We conducted an experimentation with the two specialist agents described in previous section. The DMT2 guideline (Agent D) is represented as an HTN domain with 42 primitive tasks and 13 compound tasks, and the AH guideline (Agent A) includes 42 primitive tasks and 22 compound tasks.

Figure 16 depicts multiple scenarios, each one representing the health conditions for one same patient at different time points. Our purpose is to show that the architecture is able to tackle different patient states along time (as a consequence of the patient response to the treatment over time). For the sake of clarity, we intentionally obviate the *diagnosis* subgoal in this validation and focus instead exclusively on *treatment* plans. For each scenario, the second column textually describes the findings of a clinician for the given patient (anamnesis). The patient profile is represented by the health conditions and the temporal information that embodies the currently applied therapy to the patient. It is worth observing that since health conditions are temporally annotated, the architecture is able to detect whether the patient health conditions have changed from one scenario to another. Moreover, temporal information about the current therapy which was prescribed in the previous medical consultation is represented as well. Concretely, we show the following temporal information for each scenario:

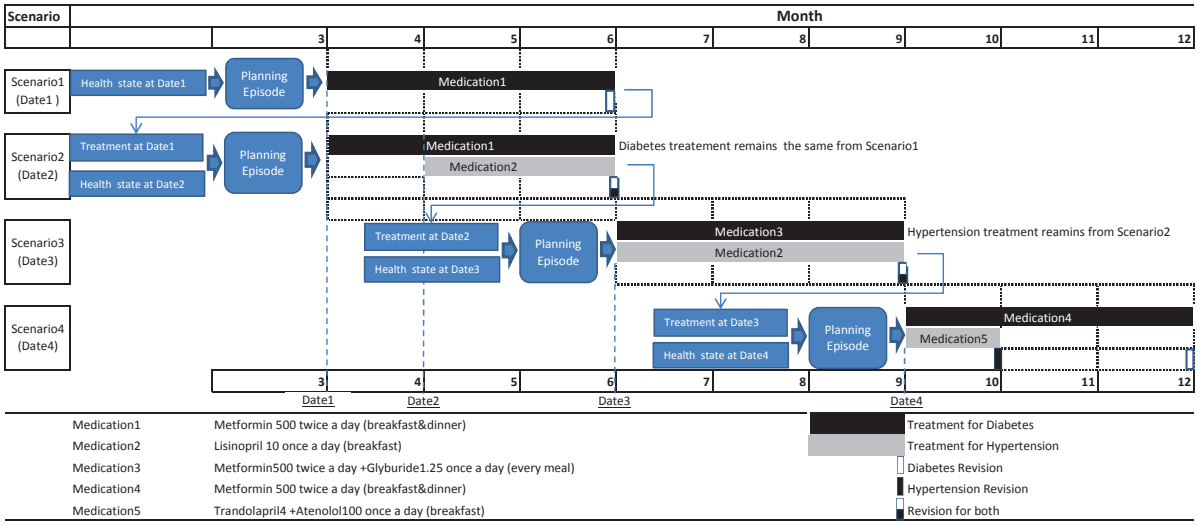

Figure 17: Conciliated plans generated for each scenario. A plan for *Scenario[i]* is generated in a planning episode at *Date[i]*, considering the patient health conditions at *Date[i]* and the treatment plan generated at *Date[i-1]*, i = 2..4.

- prescribed medication: for example, the fact *(Treatment Metformin 500 oral solution twice_a_day)* means the patient is taking Metformin 500, oral mode, format as solution and frequency twice a day when the patient comes to medical consultation; this therapy was prescribed in the previous visit
- duration of the treatment: for example, the fact *(duration_Treatment Metformin 2160hrs)* represents that the treatment prescribed in the last visit has a duration of 3 months (2160 hours)
- start time of the revision of a prescribed treatment: for example, the fact *(start_FollowUp HyperTension "28/06/2018 09:00:00")* represents that a Hypertension follow-up scheduled for "28/06/2018 09:00:00" was set in the previous medical visit

In summary, if at time t the patient is under a treatment which was prescribed at the previous scenario $t-1$, this is represented in the fourth column of the scenario at time t in Figure 16. Consequently, the prescribed therapy is taken into account when generating the conciliated treatment plan for the health conditions of the patient at time t.

The multiple scenario setting is intended to show the flexibility of the architecture to cope with situations wherein the conditions of each disease progress separately and thereby comorbidity conditions are not needed to be synchronized all at once. Concretely, *Scenario 1* represents the state of a patient at a given time with only one disease (Diabetes). *Scenario 2* represents a new state (at a later time) of the same patient, who becomes a comorbid patient because of elevated levels of blood pressure (Diabetes and Hypertension). *Scenario 3* represents a new state of the same patient in which only the conditions of one disease (Diabetes) have changed. *Scenario 4* represents a later state of the same patient in which the health conditions of both diseases have changed concurrently.

Figure 17 shows for each scenario several conciliated treatment plans at different planning episodes generated by the architecture. The chronological evolution of the different scenarios is as follows[6]:

[6]It is important to note that this is a simulation, and for the sake of simplicity we have represented abrupt changes in medication that might not correspond to real practice. We are aware that in real clinical cases the cessation or exchange of medication is not as abrupt as shown in this chronological sequence.

- *Scenario1*: The conciliated treatment consists of medication for only one disease (Diabetes) since the patient has not yet been diagnosed with AH. The Diabetes Guideline recommends, given the health conditions of the patient in this scenrio (A1C = 7), to prescribe Metformin during three months and a follow-up visit for revising the patient's diabetes conditions.

- *Scenario2*: The conciliated treatment generated for this case is an adaptation of the treatment generated in Scenario1. The adapted treatment incorporates new medication for AH (Lisinopril) given that the Hypertension conditions in this scenario (SBP=131/DBP=82) represent an elevated blood pressure (BP). The diabetes treatment prescribed in Scenario1 remains unaltered because the patient's diabetes conditions have not changed. Therefore, following the recommendations of the Hypertension Guideline, the treatment plan for Diabetes prescribed in Scenario1 is modified by adding one task that represents the prescription of Lisinopril during two months and another task that schedules a follow-up visit for 28/06/2018 (*Date3* in Figure 17) to check patient's AH conditions. Note that, in this case, both revisions of Diabetes and AH are scheduled for the same date. That is, the planner adapts the follow-up date of AH to the date of Diabetes.

- *Scenario3* is a symmetric case to Scenario2, wherein Diabetes conditions have changed and a new treatment must be generated for this disease but the AH treatment remains the same, except that a new follow-up date is scheduled for AH. The diabetes Guideline prescribes a combined therapy of two drugs (Metformin500+ Glyburide1.25) when the glycated hemoglobin (A1C) is beyond 7.5 during three months. Just as in *Scenario2* the new treatment plan also includes a follow-up session of both Diabetes and Hypertension for 30/09/2018 (*Date4* in Figure 17).

- *Scenario4* assumes that three months have passed since the last visit and that the patient health conditions have changed for both diseases. In this case the conditions of Diabetes have improved and those of AH have worsened. In both cases, it is necessary to generate a new conciliated plan that incorporates treatments to respond to the new patient condition. Since the Diabetes conditions have improved, the medication for Diabetes is again the same as the one prescribed in Scenario1. However, the treatment for AH is intensified with a combination of two drugs (Trandolapril4+Atenolol100). In addition, a follow-up session of AH is scheduled for the following month, and another one for Diabetes within three-month time.

As it can be seen, our approach is able to cope with several cases in which the status of a patient evolves over time, and the therapy is decided on the basis of such evolution; that is, the therapy is determined at the different patient consultation dates, choosing each time a treatment adapted to the patient response to the previous treatment. Moreover, the treatment is scheduled according to the current date of the visit as well as to the recommendations of each clinical guideline. This experimentation shows that the planning activity is fully adapted to the requirements of the guidelines, and that the planning horizon (i.e. the point where the planning activity reaches to) depends on the frequency of visits (follow-up sessions) recommended by the guidelines.

In Scenario2 both agents are in charge of elaborating a *treatment* for the patient (same subgoal of the Common goal) but only the Hypertension agent calculates a plan since the Diabetes agent detects that patient's Diabetes conditions are unchanged (as explained in Section 5.5). This highlights that although the two agents are dealing with the same treatment subgoal in each scenario, each agent interprets the goal differently and accomplishes it accordingly to its own knowledge (encoded in the domain) and the temporal information encoded in the initial state.

The different plans are generated according to the clinical guidelines, but also taking into account the patient preferences, as described above. Except for the first scenario, in which only one agent participates, the application of

		PROFILE AND PREFERENCES			CONCILIATED PLAN					
Scenario	**Ag.**	**Renal**	**Side pref.**	**Intake pref.**	**Medication**	**Dosage**	**Frequency**	**Intake**	**Duration**	**Rev. Date**
Scenario1 (30/03/2018)	D	No	hypo over gi	breakfast	**Metformin**	**500**	**twice a day**	**breakfast-dinner**	**3 months**	**28/06/2018**
Scen1-WhatIf-1	D	Yes	hypo over gi	breakfast	Miglitol	25	every meal	breakfast-lunch-dinner	3 months	28/06/2018
Scen1-WhatIf-2	D	Yes	gi over hypo	breakfast	Glyburide	1.25	once a day	breakfast	3 months	28/06/2018
The Diabetes Guideline recommends Metformin as first medication unless it is contraindicated because adverse effects when the patient suffers kidney disease. In this case, Alpha Glucosidase Inhibitors (as Miglitol) are recommended. If additionally the patient prefers to avoid gastro-intestinal risks, a drug of class Sulfonylureas is recommended (Glyburide). Intake time preferences cannot be fulfilled in plans 1 and 2, since the only drug administration for which the daily frequency is once a day (an the intake preference for breakfast can be fulfilled) is Glyburide, other drugs cannot be administered as preferred by the patient.										
Scenario2 (26/04/2018)	D	No	hypo over gi	breakfast	**Metformin** **Lisinopril**	**500** **10**	**twice a day** **once a day**	**breakfast-dinner** **breakfast**	**3 months** **2 months**	**28/06/2018** **28/06/2018**
	A	No	hypo over gi	breakfast	Metformin Lisinopril	500 10	twice a day once a day	breakfast-dinner breakfast	3 months 1 month	28/06/2018 29/05/2018
Scenario2 WhatIf	D	Yes	hypo over gi	breakfast	Miglitol Chlortalidone	25 25	three a day once a day	every meal breakfast	3 months 1 month	28/07/2018 29/05/2018
	A	Yes	hypo over gi	breakfast	Miglitol Chlortalidone	25 25	three a day once a day	every meal breakfast	3 months 1 month	28/07/2018 29/05/2018
At Scenario2, Agent D proposes to continue with Diabetes therapy, a duration of 2 months for the treatment of Hypertension, and schedules the revision of Hypertension at the same date that revision of Diabetes. This is due to the resolution of a redundant intervention, as one of the many illustrated in Section 7.1. Agent A gives priority to the Hypertension Guideline that recommends a follow-up session after one month. At Scenario2-WhatIf, none agent recommend to continue with the current Diabetes treatment, since the additional qualitative preference "renal" (i.e. the patient prefers avoid the risk of kidney adverse effects) force to choose different drugs. Both agents recommend the same treatments because patient preferences hardly limit the possible choices of treatments.										

Table 7: Several plans generated for different combinations of qualitative preferences.

the life cycle of the architecture follows the same process as in the example presented in Section 7. Particularly, two plans are generated, each one by a different agent, and are evaluated according to the quantitative preferences. Clinical guidelines actually offer a large number of alternatives for medication and the main knowledge source to decide the medication to prescribe, besides the clinical knowledge which is the top-priority element, are the patient qualitative preferences.

Table 7 shows the qualitative preferences of the conciliated plans of Scenario1 and Scenario2 of Figure 17 as well as the output of the architecture in case that other different preferences of the patient were considered (*What-if* episodes). Clinicians have positively evaluated the display of alternative episodes since they are very useful to inform the patients about multiple plans with different medication prescriptions. Thereby, patients can make a better informed decision considering their own preferences.

In the *What-if* planning episodes of Table 7, the profile exhibits the patient has renal problems. This information is decisive for the clinician to determine the most adequate drug family according to the clinical guideline. Additionally, column *[Side pref]* shows qualitative preferences related to medication side effects (as the ones explained in Section 4.1 as *Side effects* vs. *Side effects*). For example, [*hypo over gi*] stands for "the patient prefers a medication that reduces risk of hypoglycemia over another that has less risk of gastro-intestinal problems", or viceversa, [*gi over hypo*]. On the other hand, qualitative preferences related to administration attributes (as those explained in Section 4.1 as *Oral medication* vs. *Injectable medication*, or *Once a day* vs. *Twice a day* vs. *Each meal*), are shown in column *[Intake pref.]*, which only shows the preferred time to take drugs.

Plans highlighted in bold face are the final conciliated plans of Scenario1 and Scenario2 of Figure 17 that result from the evaluation of the patient quantitative preferences. The *What-if* episodes show conciliated plans generated by each agent with additional preferences. For each planning episode, a comment about the main decisions made by the planning process is shown. As an example, the first plan of Scenario1 in Table 7 shows a plan generated for a patient without renal problems, who prefers avoiding the risk of hypoglycemia and taking medication once a day during breakfast. In the third plan of this scenario, the patient has renal problems and would rather avoid gastro-intestinal

adverse effects as well as taking medication once a day during breakfast.

7.4. Discussion

The final conciliated plans obtained in the validation were shown to the clinicians. The results in Table 7 show several what-if episodes that were used by the clinicians to simulate different scenarios and assess to what extent the plans generated by our approach and the plans they would have elaborated are alike. Clinicians validated the correctness of the conciliated treatment plans and valued very positively the capability of the architecture to assimilate different qualitative preferences in the design of a treatment plan for a comorbid patient based on the conciliation of several guidelines. They also stressed the value of the quantitative preferences to allow the patient select one treatment over the others, as shown in the example of Table 6. In general, clinicians enthusiastically encourage the use of patient preferences in their daily clinical practice as long as the patient has been correctly informed about the consequences of the decision.

Interestingly, the use of quantitative preferences is a distinguishable feature of our approach with respect to the work presented in [34], which also deals with preferences of multimorbid patients. One can easily associate the patient preferences and mitigation process of [34] with our qualitative preferences and preference interaction detection procedure, respectively. However, quantitative preferences, as personal options of the patient that do not entail adverse clinical interactions but simply represent personal choices, are neglected in [34]. For instance, the treatment duration or the number of medical visits are not determinant factors for a clinically valid treatment; they simply denote aspects that will make the patient life more comfortable. On the other hand, the inherently distributed computation of our agent-based approach promotes *diversity* by suggesting a series of different treatment plans that encompass the knowledge of physicians (agents) and enable patients to express their preferable choices.

One limitation argued by clinicians is that our approach does not support plan execution, patient monitoring or clinician-patient interaction during the planning process. We intend to overcome these issues in future development as described in the next section.

8. Open issues

Our tool relies upon well-known computational formalisms such as existing planning technology to reason about interactions or standard agent communication protocols to simulate exchange of knowledge amongst physicians. Hence, the model is flexible enough so as to be able to plug in other modules that enable plan execution or computer-physician interaction. Currently, our tool is aimed at calculating a plan but our intention is to leverage the experience of some of the authors of this paper in the identification of the characteristics of patient-tailored treatment plans execution on single-agent planning [27, 36]. Concretely, it is worth noting the need of adaptation to deviations like the evolving patient conditions or the high dynamism and unpredictability of the clinical environment. The execution of a multi-agent conciliated plan, on the other hand, can be envisioned with a dedicated execution agent per *Specialist* agent responsible of monitoring the actions resulting from their decision procedures (distributed execution); or with a single agent executing all of the plan actions (centralized execution). Distributed execution is commonly adopted in domains where control of plan entities is costly [45]. In the particular case of a comorbid setting, even though knowledge and planning abilities are distributed across specialist agents, every agent is able to build a conciliated plan using the decision procedures of the other agents and so monitoring and executing a plan for a comorbid patient can be done by a single execution agent. Our conciliated plan approach thus enables a centralized execution, which is more flexible and effective for both the patient and institution because it avoids the need of communication and interaction

among agents. Moreover, a centralized execution facilitates the integration with standard tools for process execution like BPM[7] runtime engines and consoles [46]. These tools provide support for user interaction, ubiquitous access and collaborative execution of a treatment plan, allowing for a rapid prototyping methodology and an agile capture of user requirements [36]. Nevertheless, when these tools are integrated with planning technologies, appropriate techniques capable of detecting and handling medical exceptions need to be incorporated (as we have shown in [36]). Regarding our MAP approach, in case of a deviation that compels calculating a new plan, all planning agents will be likewise involved in the planning process independently of the execution modality.

Another lacking feature of our approach is the capacity of computer-physician interaction. A mixed-initiative approach would enable physicians to interactively intervene during the construction of the plan, to guide the planner to adjust the plan to the patient health status or to track the plan execution. The fact that the conciliation process is orchestrated by the Initiator agent opens up the way to the participation of clinicians and patients in a mixed-initiative process, making each step of a plan selection be a decision shared by clinician and patient with the help of the scores and alternative plans provided by the architecture. Moreover, our approach highlights other features that pave the way towards a mixed-initiative tool. Most importantly, it draws upon on knowledge-driven process that takes into account the expertise, knowledge and opinions of the clinicians when designing the decision procedures that encode the clinical recommendations of the CIGs represented as HTN domains. In real scenarios, clinicians would hardly agree to execute a plan without prior consideration of the decisions made to build a plan, and this is the reason they were deeply involved in the design of the decision procedures and patient's preferences to guide the automated planning process. All in all our approach would never be categorized as a black-box system because an HTN planner is based on a deliberative planning process guided by human expert knowledge. Not only all the decisions made by an HTN planner are guided by human knowledge (previously encoded by knowledge engineers in collaboration with clinicians), but all these decisions along with the reasoning process keep recorded in an internal structure named decomposition tree, and with the appropriated machinery to provide user-friendly explanations (this is a matter of Human Computer Interaction and falls out of the scope of this paper) all these decisions can be communicated to the physician.

Another limitation is related to the burden of knowledge engineers to operationalize clinical knowledge into our architecture. Automated knowledge acquisition and model elicitation actually present a bottleneck that negatively affects the generalized adoption of knowledge-based clinical technologies. In fact, the development of methodologies and tools to translate CPGs into Computer Interpretable Guidelines (CIGs) is still an active research area and various alternatives to knowledge acquisition and authoring tools have been proposed in the literature [6]. In section 3 we argued that HPDL can be regarded as an instrumental language for making clinical knowledge actionable. HPDL is a domain-independent planning language that has been successfully applied to several domains [47, 48, 49, 12]. However, it presents some limitations for knowledge elicitation mainly because it was developed as a generic language and not particularly oriented to clinical knowledge representation. On the other hand, we have shown in this paper that HPDL offers sufficient expressiveness and flexibility to represent and reason about not only a formalized clinical guideline itself, but also patient profiles, possible resources, or other potentially relevant requirements or constraints (e.g., drug costs or staff availability). Other CIG languages like Asbru [37], PROforma [20] or GLARE [50] are specifically designed for clinical domains, and they provide their own methodologies for clinical knowledge acquisition, elicitation and representation. The translation scheme of Asbru-specified CIGs into our HPDL language proposed in [12] can be integrated with virtually any off-the-shelf methodology for medical knowledge elicitation, thus revealing

[7] Business Process Management

that state-of-the-art authoring tools and methodologies are also applicable to our approach. Hence, it is reasonable to assume that developing a system upon our MAP approach would require a similar effort as with any equivalent knowledge representation scheme.

Additionally, a key step to augment portability of the system is concerned with its integration with EHR data. Since the data model of HPDL is based on a hierarchy of types along a set of predicates and functions (with typed parameters referring to the types hierarchy), it is straightforward to come up with a systematic procedure that maps classes of an ontology into the type hierarchy, and properties and relations into functions and predicates, respectively. Moreover, our research group has developed a graphical editing tool of HPDL language based on an UML model [28] that leverages automated integration with external databases, thus minimizing the effort of integrating and adapting the system to different scenarios.

9. Conclusions

This paper presents a Multi-Agent Planning proposal to support the conciliation of clinical guidelines in comorbidity and handles two types of patient preferences. Each agent encapsulates the knowledge and expertise of a physician on a particular disease and embodies a single-disease CIG formalized as an HTN domain with activities, goals and patient preferences. The hierarchical representation enables the definition of compound activities around different decomposition methods, which are subject to the satisfaction of a medical condition or a patient preference.

Qualitative preferences activate conditional decomposition methods which determine the activities or drugs that best suit the patient desires. When conciliation of drugs or activities is not doable, this preference interaction is undertaken by selecting an alternative method of a task. On the other hand, quantitative preferences capture the troublesome guidelines of a treatment to a patient and are used to select a treatment from many possible options. Hence, our approach handles the most relevant interactions identified by clinicians as well as the interactions that emerge when dealing with *personalized* CIGs.

The multi-agent architecture revolves around two concepts: *coordination* to attain the domain merging or conciliation of CIGs, and *planning* to synthesize a multi-disease personalized and conciliated treatment (plan). At coordination, agents exchange their preference-filtered decision procedures and create a personalized merged CIG that conciliates the clinical recommendations of all the CIGs. At planning time, agents generate a conflict-free conciliated treatment plan solving the clinical and preference interactions, giving priority to their own decisions in case of conflict or requesting an alternative decomposition method in case of a preference interaction. Since each agent will most likely end up with a different conciliated plan, the patient quantitative preferences are applied through a MDCM process to select the most patient-tailored plan. This process also accounts for medical institution preferences. The evaluation function of the MCDM process is used to resolve equipoise situations and can be seen as an agreement procedure between agents in order to decide the best conciliated plan, which is later assimilated by every agent in order to proceed with the next subgoal of the common goal.

The results of the case study show that MAP is a flexible technology to support clinical decision making in comorbidity considering patient preferences. Extending single-agent planning for clinical decision support towards automated MAP technology for personalized conciliation of multiple guidelines requires: (1) providing additional knowledge to represent sources of potential interactions as well as patient preferences; (2) providing agents with a coordination mechanism to conciliate guidelines by detecting and managing interactions between CIGs; and (3) providing a mechanism to resolve equipoise situations.

The MAP architecture brings several advantages that are worth mentioning:

- **Flexibility**: we propose to dynamically exchange portions of HTN domains (decision procedures) instead of fully-instantiated plans as in classical multi-agent plan merging. That is, agents exchange the decision logic to build plans, what leads to a more flexible interaction solving process. Unlike plan merging, in our approach agents backtrack over the HTN domain representation sent by the other agents when interactions are found so they are not compelled to request a new plan and restart the whole process.

- **Distributed Knowledge Representation**: the novel proposed knowledge representation (CIGs formalized as HTN planning domains), together with the dynamic exchange of knowledge among agents, allows a *Specialist* agent to gather the best clinical evidence of other CIGs for which it is not a specialist.

- **Potentially scalable**: the encapsulation of knowledge in each agent leads to a loosely-coupled architecture that opens the way to a scalable management of new diseases. A new disease might be incorporated into the life cycle with a minimal impact in the performance of the system. Given a single-disease CIG represented as an HTN domain, the only requirement to include a new disease is to extend the knowledge representation as illustrated in the case study. Hence, providing a new agent with a CIG represented in such terms should suffice for the management of a comorbid patient with more than two diseases.

An interesting future work is to exploit the MAP step-wise process to give support to *Shared Decision Making*. The architecture should be regarded as a core technology to provide support to both clinicians and patients when making decisions before the enactment of the conciliated plan. This way, agents representing other health-care stakeholders, such as patients, evaluation units or medical institutions, can be easily incorporated in the architecture. The step-wise cycle would enable agents to communicate their local solutions either to clinicians, patients or even institution managers, who can confirm whether the proposed plans fit their preferences, or even make alternative proposals. The architecture can thus be used for what-if analysis by both institution managers and specialists to propose different treatments of a patient (that might include aspects like material and human resources).

References

[1] C. M. Boyd, M. Fortin, Future of multimorbidity research: How should understanding of multimorbidity inform health system design, Public Health Rev 32 (2) (2010) 451–474.

[2] J. A. López-Vallverdú, D. Riaño, A. Collado, Rule-based combination of comorbid treatments for chronic diseases applied to hypertension, diabetes mellitus and heart failure, in: BPM 2012 Joint Workshop, ProHealth 2012/KR4HC 2012, Revised Selected Papers, 2012, pp. 30–41.

[3] M. Dawes, Co-morbidity: we need a guideline for each patient not a guideline for each disease, Family Practice: the International Journal for Research in Primary Care 27 (1) (2010) 1–2.

[4] L. Piovesan, G. Molino, P. Terenziani, An ontological knowledge and multiple abstraction level decision support system in healthcare, Decision Analytics 1 (8) (2014) 1–24.

[5] G. García Sáez, M. Rigla Cros, I. Martínez Sarriegui, E. Shalom, M. Peleg, T. Broens, B. Pons, E. Caballero Ruiz, E. J. Gómez Aguilera, M. E. Hernando Pérez, Patient-oriented computerized clinical guidelines for mobile decision support in gestational diabetes, J Diabetes Sci Technol 8 (2) (2014) 238–246.

[6] M. Peleg, Computer-interpretable clinical guidelines: A methodological review, Journal of Biomedical Informatics 46 (4) (2013) 744–763.

[7] M. Peleg, Y. Shahar, S. Quaglini, A. Fux, G. García-Sáez, A. Goldstein, M. E. Hernando, D. Klimov, I. Martínez-Sarriegui, C. Napolitano, E. Parimbelli, M. Rigla, L. Sacchi, E. Shalom, P. Soffer, Mobiguide: a personalized and patient-centric decision-support system and its evaluation in the atrial fibrillation and gestational diabetes domains, User Model. User-Adapt. Interact. 27 (2) (2017) 159–213.

[8] L. A. Lenert, D. J. Cher, M. K. Goldstein, M. R. Bergen, A. Garber, The effect of search procedures on utility elicitations, Medical Decision Making 18 (1) (1998) 76–83.

[9] P. F. Brennan, Improving health care by understanding patient preferences, Journal of the American Medical Informatics Association 5 (3) (1998) 257–262.

[10] C. M. Boyd, J. Darer, C. Boult, L. P. Fried, L. Boult, A. W. Wu, Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases: Implications for pay for performance, JAMA 294 (10) (2005) 716–724.

[11] I. Sánchez-Garzón, J. Fdez-Olivares, E. Onaindía, G. Milla, J. Jordán, P. Castejón, A multi-agent planning approach for the generation of personalized treatment plans of comorbid patients, in: Artificial Intelligence in Medicine, AIME, Springer, 2013, pp. 23–27.

[12] A. González-Ferrer, A. ten Teije, J. Fdez-Olivares, K. Milian, Automated generation of patient-tailored electronic care pathways by translating computer-interpretable guidelines into hierarchical task networks, Artificial Intelligence In Medicine 57 (2013) 91–109.

[13] D. Riaño, A. Collado, Model-based combination of treatments for the management of chronic comorbid patients, in: Artificial Intelligence in Medicine, AIME, Springer, 2013, pp. 11–16.

[14] B. Jafarpour, S. S. R. Abidi, Merging disease-specific clinical guidelines to handle comorbidities in a clinical decision support setting, in: Artificial Intelligence in Medicine, AIME2013, Springer, 2013, pp. 28–32.

[15] S. Wilk, M. Michalowski, X. Tan, W. Michalowski, Using first-order logic to represent clinical practice guidelines and to mitigate adverse interactions, in: Knowledge Representation for Health Care - 6th International Workshop, KR4HC 2014, 2014, pp. 45–61.

[16] V. Zamborlini, R. Hoekstra, M. D. Silveira, C. Pruski, A. ten Teije, F. van Harmelen, A conceptual model for detecting interactions among medical recommendations in clinical guidelines - A case-study on multimorbidity, in: Knowledge Engineering and Knowledge Management - 19th International Conference, EKAW 2014, 2014, pp. 591–606.

[17] L. Piovesan, P. Terenziani, A mixed-initiative approach to the conciliation of clinical guidelines for comorbid patients, in: Knowledge Representation for Health Care - AIME 2015 International Joint Workshop, KR4HC/ProHealth 2015, 2015, pp. 95–108.

[18] V. Zamborlini, M. D. Silveira, C. Pruski, A. ten Teije, E. Gelejin, M. van der Leeden, M. Stuiver, F. van Harmelen, Analyzing interactions on combining multiple clinical guideline, Artificial Intelligence in Medicine 81 (2017) 78–93.

[19] E. Lozano, M. Marcos, B. Martínez-Salvador, A. Alonso, J. Alonso, Experiences in the development of electronic care plans for the management of comorbidities, in: Knowledge Representation for HealthCare, KR4HC, Vol. 5943, 2010, pp. 113–123.

[20] D. R. Sutton, J. Fox, The syntax and semantics of the PROforma guideline modeling language, Journal of the American Medical Informatics Association 10 (5) (2003) 433–443.

[21] S. Abidi, A conceptual framework for ontology based automating and merging of clinical pathways of comorbidities, in: Knowledge Management for Health Care Procedures, K4CARE, Vol. 5626, 2009, pp. 55–66.

[22] S. Wilk, W. Michalowski, M. Michalowski, K. Farion, M. M. Hing, S. Mohapatra, Mitigation of adverse interactions in pairs of clinical practice guidelines using constraint logic programming, Journal of Biomedical Informatics 46 (2013) 341–353.

[23] M. Michalowski, S. Wilk, W. Michalowski, D. Lin, K. Farion, S. Mohapatra, Using constraint logic programming to implement iterative actions and numerical measures during mitigation of concurrently applied clinical practice guidelines, in: Artificial Intelligence in Medicine, AIME, Springer, 2013, pp. 17–22.

[24] GLINDA: GuideLine INteraction Detection Architecture, http://glinda-project.stanford.edu/.
URL `http://glinda-project.stanford.edu/`

[25] D. Isern, A. Moreno, D. Sánchez, Á. Hajnal, G. Pedone, L. Z. Varga, Agent-based execution of personalised home care treatments, Applied Intelligence 34 (2) (2011) 155–180.

[26] J. Fdez-Olivares, L. Castillo, J. A. Cózar, O. García Pérez, Supporting clinical processes and decisions by hierarchical planning and scheduling, Computational Intelligence 27 (2011) 103–122.

[27] J. Fdez-Olivares, I. Sánchez-Garzón, A. González-Ferrer, J. A. Cózar, A. Fdez-Teijeiro, M. R. Cabello, L. Castillo, Task network based modeling, dynamic generation and adaptive execution of patient-tailored treatment plans based on smart process management technologies, in: Knowledge Representation for HealthCare, KR4HC, Vol. 6924, 2012, pp. 37–50.

[28] F. Palao, J. Fdez-Olivares, L. Castillo, O. García, An extended HTN knowledge representation based on a graphical notation, in: Workshop on Knowledge Engineering for Planning and Scheduling (KEPS), 2011, pp. 126–135.

[29] L. Castillo, K. Swenson, S. Reddy, L. Fischer, Clinical Decision Support System for Oncology using Adaptive Case Management: True dynamic and knowledge-based treatment process by Cognocare, Future Strategies Inc, 2014.

[30] G. Milla-Millán, J. Fdez-Olivares, I. Sánchez-Garzón, D. Prior, L. Castillo, Knowledge-driven adaptive execution of care pathways based on continuous planning techniques, in: Process Support and Knowledge Representation in Health Care, Springer, 2013, pp. 42–55.

[31] Y. Dimopoulos, M. A. Hashmi, P. Moraitis, μ-satplan: Multi-agent planning as satisfiability, Knowledge-Based Systems 29 (2012) 54–62.

[32] V. R. Lesser, K. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, D. E. Neiman, R. M. Podorozhny, M. V. N. Prasad, A. Raja, R. Vincent, P. Xuan, X. Zhang, Evolution of the GPGP/TÆMS domain-independent coordination framework, Autonomous Agents and Multi-Agent Systems 9 (2004) 87–143.

[33] A. Torreño, E. Onaindia, O. Sapena, FMAP: distributed cooperative multi-agent planning, Appl. Intell. 41 (2) (2014) 606–626.

[34] M. Michalowski, S. Wilk, D. Rosu, M. Kezadri, W. Michalowski, M. Carrier, Expanding a first-order logic mitigation framework to handle multimorbid patient preferences, in: AMIA 2015, American Medical Informatics Association Annual Symposium, 2015.

[35] L. Castillo, J. Fdez-Olivares, O. García-Pérez, F. Palao, Efficiently handling temporal knowledge in an htn planner., in: International Conference on Automated Planning and Scheduling (ICAPS), 2006, pp. 63–72.

[36] I. Sánchez-Garzón, A. González-Ferrer, J. Fernández-Olivares, A knowledge-based architecture for the management of patient-focused care pathways, Applied intelligence 40 (3) (2014) 497–524.

[37] Y. Shahar, S. Miksch, P. Johnson, The asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines, Artificial intelligence in medicine 14 (1) (1998) 29–51.

[38] P. Terenziani, E. German, Y. Shahar, The temporal aspects of clinical guidelines., Studies in Health Technology and Informatics 139 (2007) 81–100.

[39] P. Votruba, S. Miksch, A. Seyfang, R. Kosara, Tracing the formalization steps of textual guidelines, Studies in health technology and informatics 101 (2004) 172–176.

[40] M. Peleg, S. Tu, J. Bury, P. Ciccarese, J. Fox, R. A. Greenes, R. Hall, P. D. Johnson, N. Jones, A. Kumar, et al., Comparing computer-interpretable guideline models: a case-study approach, Journal of the American Medical Informatics Association 10 (1) (2003) 52–68.

[41] A. K. Mackworth, Consistency in networks of relations, Artificial Intelligence 8 (1) (1977) 99 – 118. doi:https://doi.org/10.1016/0004-3702(77)90007-8.
URL http://www.sciencedirect.com/science/article/pii/0004370277900078

[42] E. Triantaphyllou, Multi-Criteria Decision Making Methods: A Comparative Study, Vol. 44 of Applied Optimization Series, Kluwer Academic Publishers, 2000.

[43] A. J. Garber, Consensus statement by the american association of clinical endocrinologists and american college of endocrinology on the comprehensive type 2 diabetes management algorithm 2018 executive summary, Endocrine Practice 24 (1) (2018) 91–120.

[44] P. K. Whelton, 2017 ACC/AHA/AAPA/ABC/ACPM/AGS/APhA/ASH/ASPC/NMA/PCNA Guideline for the Prevention, Detection, Evaluation, and Management of High Blood Pressure in Adults, Hypertension.

[45] K. L. Myers, CPEF: A continuous planning and execution framework, AI Magazine 20 (4) (1999) 63–69.
URL http://www.aaai.org/ojs/index.php/aimagazine/article/view/1480

[46] U. Dayal, M. Hsu, R. Ladin, Business process coordination: State of the art, trends, and open issues., in: VLDB, Vol. 1, 2001, pp. 3–13.

[47] J. Fdez-Olivares, L. Castillo, O. Garcıa-Pérez, F. Palao, Bringing users and planning technology together. experiences in siadex, in: Proc ICAPS, 2006, pp. 11–20.

[48] Samap: An user-oriented adaptive system for planning tourist visits, Expert Systems with Applications 34 (2) (2008) 1318 – 1332.

[49] A. González-Ferrer, J. Fernández-Olivares, L. Castillo, From business process models to hierarchical task network planning domains, The Knowledge Engineering Review 28 (2) (2013) 175–193.

[50] P. Terenziani, S. Montani, A. Bottrighi, M. Torchio, G. Molino, G. Correndo, The GLARE approach to clinical guidelines: main features, Studies in Health Technology and Informatics (2004) 162–166.