A Multi-agent Planning Approach for the Generation of Personalized Treatment Plans of Comorbid Patients*

Inmaculada Sánchez-Garzón[1], Juan Fdez-Olivares[1], Eva Onaindía[2], Gonzalo Milla[1], Jaume Jordán[2], and Pablo Castejón[2]

[1] Department of Computer Science and AI, Universidad de Granada, Spain
[2] Department of Computer Science, Universitat Politècnica de València, Spain

Abstract. This work addresses the generation of a personalized treatment plan from multiple clinical guidelines, for a patient with multiple diseases (comorbid patient), as a multi-agent cooperative planning process that provides support to collaborative medical decision-making. The proposal is based on a multi-agent planning architecture in which each agent is capable of (1) planning a personalized treatment from a temporal Hierarchical Task Network (HTN) representation of a single-disease guideline, and (2) coordinating with other planning agents by both sharing disease specific knowledge, and resolving the eventual conflicts that may arise when conciliating different guidelines by merging single-disease treatment plans. The architecture follows a life cycle that starting from a common specification of the main high-level steps of a treatment for a given comorbid patient, results in a detailed treatment plan without harmful interactions among the single-disease personalized treatments.

Keywords: multi-agent planning, comorbidity, guideline conciliation.

1 Introduction

The generation of personalized treatment plans for patients with a single disease is a widely studied problem in the field of AI in Medicine. Most of the proposed approaches implement sophisticated problem solving processes that (1) take as input a knowledge base containing a Computer Interpretable Guideline (CIG[1]); and (2) adapt the activities of the formal guideline, from patient-focused clinical data, in order to generate a personalized treatment to be performed at the point of care. Although these approaches provide decision support in order for the clinical staff to efficiently follow the guidance provided by CPGs, they either have neglected or do not have deeply addressed the problem of planning personalized treatments for comorbid patients (i.e., patients with multiple and simultaneous diseases [3]). This is mainly due to the fact that a CPG gathers

* Work partially supported by projects P08-TIC-3572 and TIN2011-27652-C03-03/01.
[1] Clinical Practice Guideline (CPG) represented in a formal language.

N. Peek, R. Marín Morales, and M. Peleg (Eds.): AIME 2013, LNAI 7885, pp. 23–27, 2013.

the clinical evidence for a concrete disease, and rarely considers patients with multiple diseases in clinical trials [8].

However, the problem of tailoring a treatment plan for a comorbid patient requires to conciliate the activities recommended by several single-disease CPGs. In this scenario, new problem solving capabilities are needed in order to detect and resolve eventual interactions (or conflicts) between different clinical knowledge sources. These interactions may be of different types, like conflicts between recommended medications of different diseases (drug-drug interactions), adverse effects on other diseases when treating the target disease (drug-disease interactions), inefficient scheduling of activities or even contradictory/redundant recommendations [3]. Regarding this problem, different perspectives have been proposed as [7], [2] or [6]. However, the use of multi-agent techniques appears to be a very suitable approach, since agents can encapsulate the expertise and skills of a specialist (or team of specialists). The strength of these techniques has been exploited in projects like GLINDA [1].

Some other features can be identified in this new scenario: (1) The medical expertise is distributed among different CPGs; (2) There is a common task that consists of the design of a personalized treatment plan, without harmful interactions, for a comorbid patient; (3) It has to be accomplished by at least two clinicians or by one clinician consulting at least two different knowledge sources and; (4) Each specialist has both disease-specific problem solving and cooperation capabilities that can exploit to carry out a collaborative medical decision-making process. These features fulfill the main requirements of a cooperative distributed problem solving process. Moreover, the generation of a personalized treatment plan from multiple CPGs, for a comorbid patient, can be addressed as a multi-agent cooperative planning process. Concretely this work presents a multi-agent planning architecture in which each agent is capable of (1) *planning* a personalized treatment for a given patient, from a temporal Hierarchical Task Network (HTN) representation of a single-disease guideline, and (2) *coordinating* with other planning agents by both sharing their local, disease-specific recommendations, and resolving the eventual conflicts that may arise when conciliating their single-disease plans. Thus, the architecture follows a life cycle that starting from a common specification of the main high-level steps of the treatment for a given patient, it will result in a personalized plan without harmful interactions.

2 Multi-agent Planning Architecture

Figure 1(a) shows the proposed Multi-Agent Planning (MAP[2]) architecture that is composed of an *Initiator Agent* and several *Planning Agents*. A planning agent can be seen as the representation of a clinical specialist capable of (1) planning a personalized, single-disease care plan (*HTN planner* module); (2) coordinating with other specialists, by sharing its single-disease recommendations and experience (*Coordinator* module) and (3) detecting and resolving conflicts between its recommendations and those of others (*Conflict Solver* module).

[2] MAP is the combination of multi-agent techniques and planning techniques [9].

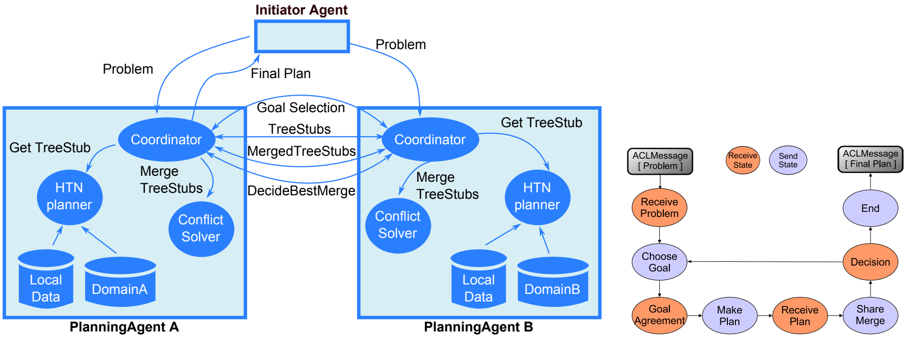

(a) Architecture diagram (b) Coordinator behaviour

Fig. 1. Multi-Agent Planning approach

The mission of the *Initiator Agent* is to send the global problem to the *Planning Agents*. A global problem may be defined by a starting time point and a list with the high-level steps of the care plan. For instance, the following code shows a problem denoting that the phases of diagnosis, treatment and follow-up have to be designed for the patient $?p$ starting from a specific date: $((\$start == "2012/12/21")(Diagnosis\ ?p)(Treatment\ ?p)(FollowUp\ ?p))$.

Note that the process of designing a treatment plan for a comorbid patient is a cooperative planning process that iterates over the high-level tasks of the global problem and for each task a joint treatment plan is obtained. Such joint plan is generated by a merging process developed by the *Conflict Solver* module in which single-disease treatments (automatically generated by *HTN planners*) are combined in order to detect and solve eventual interactions. Next subsections briefly describe the different modules of a *Planning Agent*.

2.1 HTN Planner : Single-Disease Treatment Planning

The HTN planner takes as inputs a *planning domain* (i.e., a hierarchy of tasks representing compound and primitive activities) and a *planning problem*. A planning domain encodes a single-disease CPG based on HPDL, a temporal HTN planning language as expressive as standard CIGs languages [5]. HPDL provides support to represent time-annotated guidelines and it is capable of representing decisions, different execution flows and iterative task decomposition schemas. A planning problem represents the set of goals to be achieved during the planning process (e.g., $(Treatment\ ?p)$) as well as the initial state, which is encoded in the block *Local Data* of Figure 1(a). Local Data contains information about patients (e.g., demographic and clinical data), drug medication patterns (e.g., dose, frequency, administration mode) as well as interactions (e.g., drug-drug interactions, drug-disease interactions). As result, the HTN planner[3] generates a *TreeStub* composed of a personalized, single-disease plan (with time-annotated primitive actions) and the rules of the planning domain applied to generate it.

[3] For more information about the planning process, please see our previous work [4,5].

2.2 Conflict Solver : Plan Merging Strategy

This module develops a plan merging process by combining the planning domains of two *TreeStubs*. This means that every conflict solver uses two different knowledge sources[4] for generating a merged, conflict-free treatment plan. Interactions are avoided by means of the *preconditions* and *effects* (expressed in HPDL) of the primitive actions of the planning domain. In this way, before including a new action to the joint plan, the following requirements must be satisfied: (1) this action is not contraindicated for the patient, (2) this action does not generate any drug-drug or drug-disease conflict. In planning time, such conditions are checked by comparing the clinical data of the patient (e.g., list of prescribed drugs, diagnosed diseases and contraindicated actions) with the knowledge extracted from CPGs (e.g., interactions between drug-drug and drug-disease). Note that this data are encoded in the Local Data block. Moreover, after adding an action to the plan, its effects are applied updating the clinical data of the patient.

2.3 Coordinator : Multi-agent Planning Cooperation

Figure 1(b) shows a finite state machine model of the behaviour of the *Coordinator* module. There are two different kind of states: (1) *receive*, where the agent waits for an *ACLMessage*[5]. from other agents, and (2) *send*, where the agent performs some actions and sends an *ACLMessage* to other agents.

In the first state (*Receive Problem*), the coordinator receives the global problem. Then, after choosing a specific task-goal of the problem and sending it to the other agent (*Choose Goal*), it waits for the reception of the other agent's goal (*Goal Agreement*). When both agents agree on the goal to solve, each one calls to its HTN planner to generate a local plan (*Make Plan*), as explained before. Then, the planner returns a *TreeStub*, which is send to the other agent. In the next state (*Receive Plan*), the coordinator receives the *TreeStub* from the other agent and then, a merging process is performed by the conflict solver in order to combine the just received *TreeStub* with its local plan (*Share Merge*). Then, agents share their merged treatments and select the best one as the final global plan. The coordinator also analyses if the planning process has finished (*End*), in which case it sends the final global plan in an *ACLMessage* to the initiator agent. Otherwise, it selects the next task-goal to achieve (*Choose Goal*).

Finally, the best plan is selected according to an utility function that may incorporate different clinical criteria: (1) *Efficiency* from the institutional point of view, by optimizing the used resources (e.g., lower cost of drugs) or the temporal cost (e.g., lower duration of treatment); (2) *Complexity* from the patient point of view, by quantifying the complexity of the medication regimen (as proposed in [3]) and selecting the easiest one. In fact, for elderly people, it is preferable a medication regimen with few drugs and few different dosage schedules.

[4] More clinical knowledge than the used previously by the HTN planner.

[5] A message whose content is based of the standard FIPA ACL (Agent Communication Language) `http://www.fipa.org/specs/fipa00061/`.

3 Conclusions and Future Work

The proposed multi-agent planning architecture provides support for the generation of treatment plans for comorbid patients by carrying out a cooperative planning process that, starting from a high-level description of the main steps of a care plan, results in an interaction-free, personalized treatment plan. Such plan is generated and merged from multiple clinical guidelines (represented as HTN domains), in a collaborative process based on both the sharing of local treatments (obtained by different temporal HTN planners) and the selection of the best proposal according to some clinical criteria (efficiency or complexity).

By means of a previous experimental evaluation, some *simulated* CPGs were represented and some preliminary results were obtained, which support this work. We are presently engaged in a detailed experimentation with real CPGs of the most prevalent diseases in comorbid patients, since our research group has already experience in the representation of real clinical knowledge in HPDL [4]. Furthermore, we intend to work on both a more collaborative plan merging approach with clinicians and on argumentation techniques to reach agreements.

References

1. GLINDA: GuideLine INteraction Detection Architecture, `http://glinda-project.stanford.edu/`
2. Abidi, S.R.: A conceptual framework for ontology based automating and merging of clinical pathways of comorbidities. In: Riaño, D. (ed.) K4HelP 2008. LNCS, vol. 5626, pp. 55–66. Springer, Heidelberg (2009)
3. Boyd, C.M., Darer, J., Boult, C., Fried, L.P., Boult, L., Wu, A.W.: Clinical practice guidelines and quality of care for older patients with multiple comorbid diseases: Implications for pay for performance. JAMA 294(10), 716–724 (2005)
4. Fdez-Olivares, J., Castillo, L., Cózar, J., García Pérez, O.: Supporting clinical processes and decisions by hierarchical planning and scheduling. Comput. Intell. 27(1), 103–122 (2011)
5. González-Ferrer, A., ten Teije, A., Fdez-Olivares, J., Milian, K.: Automated generation of patient-tailored electronic care pathways by translating computer-interpretable guidelines into hierarchical task networks. Artif. Intell. Med. (2012) (in Press)
6. Hing, M.M., Michalowski, M., Wilk, S., Michalowski, W., Farion, K.: Identifying inconsistencies in multiple clinical practice guidelines for a patient with co-morbidity. In: IEEE International Conference on BIBM Wkshp., pp. 447–452 (2010)
7. Lozano, E., Marcos, M., Martínez-Salvador, B., Alonso, A., Alonso, J.R.: Experiences in the development of electronic care plans for the management of comorbidities. In: Riaño, D., ten Teije, A., Miksch, S., Peleg, M. (eds.) KR4HC 2009. LNCS, vol. 5943, pp. 113–123. Springer, Heidelberg (2010)
8. Lugtenberg, M., Burgers, J.S., Clancy, C., Westert, G.P., Schneider, E.C.: Current guidelines have limited applicability to patients with comorbid conditions: A systematic analysis of evidence-based guidelines. PLoS ONE 6(10), e25987 (2011)
9. de Weerdt, M., Clement, B.: Introduction to planning in multiagent systems. Multiagent and Grid Systems 5(4), 345–355 (2009)