Run-time Support to Comorbidities in GLARE-SSCPM

Alessio Bottrighi, Luca Piovesan and Paolo Terenziani
DISIT, Universita' del Piemonte Orientale, Viale Teresa Michel 11, Alessandria, Italy

Keywords: Clinical Guidelines, Comorbid Patients, Interaction Analysis and Management.

Abstract: Comorbidities play a relevant role in healthcare, so that, in the last years, several approaches Medical Informatics and Artificial Intelligence have developed software tools to support physicians in the treatment of comorbid patients. Computer Interpretable Guidelines (CIGs) are consolidated decision support tools to help physicians, but they are devoted to provide evidence-based recommendations for one specific disease. In order to support the treatment of patient affected by multiple diseases, challenging additional problems have to be addressed, such as (i) the detection of the interactions between CIG actions, (ii) their management, and, finally, (ii) the "merge" of CIGs. Several CIG approaches have been recently extended in order to face (at least one of) such challenging problems, and one of them is GLARE (GuideLine Acquisition Representation and Execution). However, such approaches have mostly focused on the "a-priori" treatment of such problems, while addressing them "run-time" (i.e., to support physicians during the execution of the CIGs on a specific patient) involves additional challenges, and requires additional methodologies. In this paper we take advantage of previous extensions of GLARE (to cope with issues (i), (ii), (iii)), and propose a new knowledge-based, "focused" and interactive management of comorbid patients.

1 INTRODUCTION

The term *comorbidity* indicates the co-occurrence of more than one disease in a patient. They are quite frequent (an average of 25% of the population), thus constituting an important problem from different viewpoints.

Evidence-based decision making is a quite consolidated practice in healthcare, since it exploits the evidence and knowledge provided by clinical trials, and by previous experiences. One of the main methodologies to put evidence-based medicine into practice is the development of Clinical Practice Guidelines (CPGs). CPGs are defined as "*systematically developed statements to assist practitioner and patient decisions about appropriate health care in specific clinical circumstances*" (Institute of Medicine (US), 1990). Generally, CPGs are elaborated by national or international teams of specialists, and collect and organize in a textual form the knowledge available in literature to manage a specific clinical circumstance. They play a major role in modern healthcare, and thousands of CPGs have been devised in the last few years. For instance, the Guideline International Network, which groups 97 organizations from all the continents, provides a library of more than 6000 CPGs.

Additionally, in the last 30 years or so, the research in Artificial Intelligence and in Medical Informatics has shown that software tools can be designed to increase the practical impact of CPGs in healthcare. Specifically several software tools have been devised in order to acquire, represent, execute and reason with the so-called Computer-Interpretable Guidelines (CIGs henceforth; see, for example, the surveys (Peleg, 2013; Ten Teije et al., 2008)).

1.1 CIG and Comorbidities

Unfortunately, CPGs provide evidence-based information of interventions, but only on individual pathologies. The simple solution of applying multiple CPGs (one for each disease) to a patient does not work: the treatments recommended by different CPGs may interact with each other, and such interactions may be (very) dangerous for patients. The approach of considering all the possible combinations of diseases is not only difficult, but also impractical. Such considerations

Bottrighi, A., Piovesan, L. and Terenziani, P.
Run-time Support to Comorbidities in GLARE-SSCPM.
DOI: 10.5220/0007685004980505
In *Proceedings of the 12th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2019)*, pages 498-505
ISBN: 978-989-758-353-7

highlight the importance of developing methodologies to *merge* CPGs for single disease interventions to provide professionals' assistance to comorbid patients (Riaño and Collado, 2013). The development of such methods has been identified as one of the "grand challenges" for clinical decision support (Sittig et al., 2008). Since the early 2010's, the research in Computer Science has been very active in such a challenging area of research.

1.2 State of the Art

In general, the approaches devised in such an area have specialized on the treatment of two different subproblems:

(i) the detection of interactions between CIGs, and their management (i.e., how to "solve" interactions), and
(ii) the "merge" of CIGs.

Issue (i) above has been faced by relatively few approaches in the CIG literature. In particular, the approach in (Zamborlini et al., 2014) provides a knowledge-based solution. It proposes a CIG-independent conceptual model for medical actions and reasoning forms operating on it, as well as domain-independent rules to identify different types of interactions on the basis of such a knowledge. A similar approach has been pursued in the GLARE approach (see Section 2 and (Piovesan et al., 2014), (Piovesan and Terenziani, 2015)).

On the other hand, issue (ii) has been faced by several CIG approaches. It is possible to distinguish between the approaches aiming at achieving "conservative" CIGs, and those that do not. The approach in (Sánchez-Garzón et al., 2013), for instance, belongs to the latter category. It builds ad-hoc CIGs from scratch, using an agent-based approach. Agents with hierarchical planning capabilities represent experts in the treatment of specific diseases. The CIG coping with the comorbidity is obtained through the coordination of all the agents. However, the mainstream is constituted by conservative techniques, attempting to merge existing CIGs with limited changes since, in the real medical practice, physician need to follow as much as possible evidence-based recommendations, such as the ones proposed in the CPGs (and, thus, CIGs) in the literature. The approaches in such a mainstream mostly assume that the possible interactions and their managements have been defined a priori by physicians, and focuses on CIG merge only. However, quite different techniques have been proposed. For instance, in (Wilk et al., 2013), constraint logic programming (CLP) is adopted. A CLP is derived from the CIGs, the interactions and their managements, and a mitigation algorithm is proposed in order to achieve the merge. Riaño and Collado (Riaño and Collado, 2013) propose a model-based approach for the merge. They model treatments as oriented graphs composed by decisions and actions. With the help of physicians, they define a set of operators to merge decisions or actions. The combination of the original CIGs is obtained through the application of the operators. On the other hand, in GLARE, the different management options applied to independently solve the interactions are merged through a conciliation module which is based on CSP (Constraint Satisfaction Problems) methodologies (Piovesan and Terenziani, 2016).

1.3 "Run-time" Support

Some of the above approaches can be used both (i) "a-priori", to analyse interactions between CIGs or to merge them, without any reference to a specific patient, and (ii) "run-time", to support the execution of CIGs on a specific patient. However, the "run-time" application of the above methodologies involves the resolution of new problems: *when* and *on which parts of the CIGs interaction detection* has to be performed? And the *management of the interactions*? And the *merge*? Such problems are still open problems in the specialised literature, and the goal of this paper is to propose a general methodology to cope with them, thus providing physicians with an effective, *user-friendly* and *highly interactive* approach supporting physicians in the *run-time* treatment of comorbid patients. Specifically, our approach grounds on GLARE, and on its extensions (called GLARE-SSCPM) to deal with comorbidities, which are briefly resumed in Section 2). However, we emphasize that our methodology is mostly system-independent, and can be tuned in order to apply to other approaches to comorbidities in the literature.

2 BACKGROUND: GLARE AND GLARE-SSCPM

GLARE Support System for Comorbid Patient Management (GLARE-SSCPM; (Piovesan et al., 2018)) is an extension of GLARE (Terenziani et al., 2001) to support the management of comorbidities, and which takes into account both the (i) interaction detection and management, and (ii) the CIG merge.

In the following, we briefly resume such an approach, which is the basis of the support to the approach proposed in Section 3.

2.1 GLARE

GLARE (GuideLine Acquisition Representation and Execution, (Terenziani et al., 2001)) is a well-known CIG framework, designed in a long term cooperation between the University of Eastern Piedmont and the Azienda Ospedaliera San Giovanni Battista in Turin (one of the largest hospitals in Italy), started in 1997. The kernel of GLARE provides a formalism to represent CIGs, a tool to acquire them, a mechanism to execute a CIG on a specific patient. In GLARE, CIGs are modelled as hierarchical graphs, in which nodes represent actions or decisions and arcs represent the control flow relations between nodes. GLARE distinguishes between *atomic actions* (simple steps in a CIG) and *composite actions* (plans), which are defined in terms of their components (thus supporting the definition of CIGs at different *levels of abstraction*). Atomic actions can be *work* actions (a procedure which must be executed), *pharmacological* actions (a drug to be administered), *query* actions (retrieval of information from the clinical record/examinations) or *decision* actions (choice among different alternatives). In particular, GLARE distinguish among *diagnostic* and *therapeutic* decisions (see the discussion in Section 3.2).

Arcs are used to represent the control flow relations, and can be annotated with *temporal constraints*. In particular, a *sequence* arc from node N1 to N2 indicates that the action represented by N1 must terminate before the execution of N2. On the other hand, *constrained* arcs represent complex temporal relations between nodes (e.g., N2 *during* N1), and can be used to enforce concurrent execution of actions.

The kernel of GLARE consists of two main modules: the acquisition module and the execution one. The *acquisition* module proposes a user-friendly graphical interface for the acquisition of CIGs, and stores them in an internal format. The *execution* module takes in input a CIG and the clinical record of a specific patient, and supports the "execution" of the CIG on the patient. The execution module is based on the "agenda techniques" (Terenziani et al., 2001): at each time during the execution of a CIG, GLARE determines (in the agenda) the set of current actions, each one paired with a time window, indicating when the action has to be executed (minimum and maximum time) to comply the temporal constraints in the CIG. Notably, GLARE supports concurrent actions.

GLARE's architecture is open. In the latest years, several new modules and\or methodologies have been added to cope with automatic resource-based contextualization (ADAPT module, (Terenziani et al., 2004)), temporal reasoning (TR, (Anselma et al., 2006)), decision making support (DECIDE_HELP, (Montani et al., 2005)), and model-based verification (VERIFY, (Bottrighi et al., 2010)). Recently, GLARE has been extended to cope with comorbidities (see below).

2.2 GLARE-SSCPM

GLARE-SSCPM (Piovesan et al., 2018) proposes a set of user-friendly supports to the treatment of comorbidities. It is a knowledge-based approach, aimed to support step-by-step physicians in the treatment of comorbidities.

Operationally speaking, GLARE-SSCPM is based on a CIG-independent knowledge base of clinical actions, effects, and interactions, and supports three main tasks:

(1) The detection of interactions occurring between CIGs
(2) The management of the interactions
(3) The final merging of the CIGs

Since, in the real practice, interaction occur in time, all the above tasks can be achieved only if the temporal dimension is taken into account. Therefore, also (4) *Temporal Reasoning* is considered in GLARE-SSCPM:

Knowledge Base and Reasoning Supports. GLARE-SSPCM is based on a *Knowledge Manager*, i.e., a module coping with additional (CIG-independent) medical knowledge (Piovesan et al., 2014). It adopts an OWL ontological model developed in collaboration with expert physicians, using Protégé and integrating part of medical models, such as SNOMED CT and ATC. Each action in GLARE can be associated with one or more elements of the ontological model. Such a knowledge base contains both a general ontology, describing general notions such as actions, action intentions\effects, time, interactions, as well as domain-specific knowledge, such as possible interactions between specific drug types. Moreover, the Knowledge Manager module is provided with standard OWL reasoners providing inferences. Such inferential mechanisms are used to devise a tool that navigates the knowledge base and detects which are the possible interactions (if any) between actions' effects.

Interaction Detection. GLARE-SSCPM *Interaction Detection* module (Piovesan et al., 2014) provides a flexible and interactive *focusing* tool allowing physicians to navigate through the different abstraction levels in the CIGs, to identify the "relevant" actions. Once the actions of interest are identified though focusing, interaction detection is automatic: GLARE-SSCPM exploits the knowledge provided by the knowledge manager and the OWL reasoner to retrieve all the interactions between the intentions, effects and drugs prescribed (in case of pharmacological actions) of the focused actions.

Interaction Management. Once detected and analysed, interactions must be managed. Management options are local (and as small as possible) changes in the original CIGs, which make the original GIGs executable, avoiding undesirable interactions and promoting desirable ones. On the basis of the medical literature, GLARE-SSPCM propose a wide range of *general* (i.e., CIG independent and domain independent) *interaction management options* (Piovesan and Terenziani, 2015): Safe Alternative, Replanning, Temporal Avoidance, Effect Monitoring, Dosage Adjustment, Interaction Mitigation, Interaction Alignment, Intention Alignment.

GLARE-SSCPM provides a facility to instantiate each one of such options, i.e., to apply it to a specific input interaction, and to modify the CIGs accordingly. The idea is that, given a specific interaction, the user-physicians may apply one of the options, or even trying to apply more than one, in a "what-if" modality, see what the consequences on the CIG are, and finally chose an apply the preferred option in a definitive way.

Merge. Once the interactions have been identified, and managed in isolation, the union of the original CIGs with the applied managements is not yet an executable CIG, since the management options lead to changes to the original CIGs that are "locally" consistent, but possibly not consistent with each other. For these reasons, a final "merging" step is required. In GLARE-SSCPM such a step is performed by the *CIG Conciliation* module (Piovesan and Terenziani, 2016), which provides as output a "merged" CIG executable by GLARE.

Temporal Reasoning. GLARE-SSCPM provides the *Temporal module* (Piovesan et al., 2015), to cope with temporal constraints and to perform temporal reasoning. Such a module operates as a *knowledge server*: temporal problems may be demanded to the Temporal module, which provides them a solution (or report that there is no solution).

3 RUN-TIME SUPPORT TO COMORBIDITIES

3.1 Philosophy of the Approach: Focusing and Interactivity

As discussed in the Introduction, several approaches in the literature focus on the a-priori "merge" of CIGs, to avoid dangerous interactions. Such approaches usually consider whole CIGs, and mostly operate without interacting with physicians: given two or more CIGs, to provide to physicians a new CIG, avoiding dangerous interactions.

GLARE-SSCPM follows a different philosophy: it provides a highly interactive approach, in which physicians may (i) focus on specific subparts of the CIGs, (ii) analyse possible interactions and (iii) adopt GLARE–SSPCM to check the effects of applying different management options to deal with interactions (Piovesan et al., 2018).

In this paper, we propose a methodology to extend the approaches coping with CIGs and comorbidities with proper supports for "*run-time*" execution, and we follow GLARE-SSCPM "philosophy": our methodology supports "*focusing*", and highly *interactive* with physicians.

Focusing is needed because, when executing CIGs on a specific patient, physicians are not interested with the whole CIGs, but only on the subpart of them that is applicable to the given patient, given the patient status. Indeed, focusing is needed along two dimensions:

(i) The dimension of alternative paths in the CIGs
(ii) The "temporal" dimension

Dimension (i) concern the fact that real CIGs usually contain many (even hundreds) of different alternative paths, depending on the different status that the patient may assume during the CIG execution. Obviously, only the paths that are recommended (given the current status of the patient) are interesting for physicians, and thus have to be taken into account by decision support tools.

Dimension (ii) regards the fact that physicians do not usually plan patient treatments far-away in the future. They consider a limited "window" of the CIG, usually not exceeding the next decision step in the CIG. Indeed, taking "future" decisions on the basis of the current status of the patient is nearly a "bid", which is rarely performed by physicians.

Interactivity is needed, in general, because we see our approach as a support tool, which does not substitute physicians, but helps them, by providing

additional knowledge and recommendations. Specifically, in the case of co-morbidities, while we provide a fully automatic support to the detection of possible interactions between (the "focused" parts of the) CIGs, we want to be highly interactive in the selection of the management options to treat such interactions. In general, more than one option is applicable, and we do not want to impose any specific choice to physicians. On the other hand, we want to support them in such a choice, by showing them in an automatic way the consequences of choosing a given option, or another.

3.2 Scheduled and Candidate Actions

From the practical point of view, a key issue to realize the notion of "run-time" focusing is the definition of *scheduled* (CIG) actions, and of *candidate* ones. To propose such definitions, we first have to point out the different nature of diagnostic vs therapeutic decisions in CIGs. GLARE (as well as several other CIG tools), clearly distinguishes between diagnostic and therapeutic decisions. Diagnostic decisions discriminate among different diagnoses on the basis of the patient status, considering a set of parameters (e.g., blood pressure, fever, …), which vary from decision to decision. In GLARE such decisions are represented as scored or Boolean decisions. The decision criteria are described within the decision action, and are automatically evaluated by GLARE, considering the clinical record describing the status of the patient. Though in GLARE diagnostic decisions are taken in a semi-automatic way (since GLARE allows physicians to over-rule the decision taken automatically by the system, selecting diagnoses different from the ones derived from the automatic evaluation of the decision criteria on the basis of the status of the patient), such decisions are strictly related to the state of the patient (so that physicians cannot freely choose among them, independently of the patient's status). On the other hand, in therapeutic decisions physicians have to choose among different therapies that are all recommended (by the CIG) for the given category of patients. Physicians have usually the "full control" of such therapeutic decisions, since all CIG alternative treatments are usually "eligible" for patients. The choice is done considering a given set of parameters: *effectiveness, cost, side-effects, compliance, duration*. Thus, in GLARE, a therapeutic decision action is represented by a qualitative evaluation of each one of the parameters above, for each one of the alternatives. At run-time, GLARE presents such evaluations to physicians, who are completely free to choose among each one of the alternatives.

As a consequence, at any time during the execution of a CIG on a patient, we distinguish.

(1) the set of current actions
(2) the set of scheduled actions
(3) the set of candidate actions

In Figure 1 in the following, we show a simple example of GLARE CIGs, to exemplify the definitions. For the sake of generality, we consider an abstract example, instead of a concrete one. In the example, round nodes represent work and pharmacological actions, red diamond represent diagnostic decisions, green diamonds represent therapeutic decisions, and arcs represent the control flow of actions. For the sake of simplicity, we only consider sequence arcs (so that, at each time, each CIG has only a current action).

Definition. Current Actions. The action to be executed next (multiple next actions are possible, in case of concurrency)

Example. For instance, in our example, we suppose that the current action in CIG_1 is A_2.

Definition. Scheduled Actions. Besides the current action(s), the set of scheduled actions contains the set of CIG actions which, if no failure or exception arise, have necessarily to be executed next, and their time window.

Specifically, the scheduled actions are all those CIG actions that can be reached through chains of

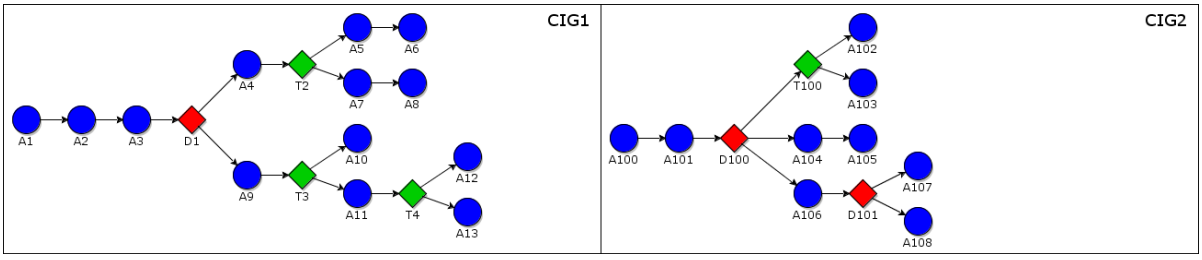

Figure 1: Example CIGs. Round blue nodes represent work and pharmacological actions, red diamond represent diagnostic decisions, green diamonds represent therapeutic decisions, and arcs represent the control flow of actions.

sequence and *constrained* arcs starting from the current actions, until a *decision* action is reached.

Example. In CIG_1, if the current action is A_2, the set of scheduled actions is $\{A_2, A_3, D_1\}$

***Defintion. Candidate Actions*.** The set of candidate actions contains the set of CIG actions which (if no failure or exception arise), can possibly be scheduled for execution, until a new therapeutic decision has to be taken by physicians.
Candidate actions include all the actions that can be reached from scheduled actions until a therapeutic decision has to be taken. The idea is that, while therapeutic decisions have to be taken by physicians, diagnostic decisions depends on the status of the patient. Therefore, the outcome of a future diagnostic decision cannot be known a-priori, and physicians may want to consider also the actions that have to be taken after such decisions.

Example. In CIG_1, if the current action is A_2, the set of candidate actions is $\{A_4, T_2, A_9, T_3\}$

3.3 GLARE Extensions

GLARE supports the execution of multiple CIGs. For each CIG to be executed on a patient, GLARE provides physicians with an ***Executor*** module supporting the execution. In our approach to run-time management of comorbidities, GLARE executor has been enriched with the possibility of sending and receiving messages to\from a Comorbidity Master Module, and to activate GLARE-SSCPM Interaction Management Module and Conciliation Module.

In the following we consider a single patient (the extension to multiple patients is trivial).

3.4 Extensions to the Executor

The modifications to GLARE's original Executor module are quite limited: it is extended to communicate with the Comorbidity Master Module. In particular, the Executor module sends to the Comorbidity Master Module a message

(i) when it is created (i.e., when the execution of a new CIG is started on the patient)
(ii) when the execution of a CIG action is (successfully) terminated. In case the action is a decision, also the selected path is sent to the Comorbidity Master Module.

It receives from the Comorbidity Master Module a message whenever

(iii) one or more interactions have to be managed

When the Executor receives a message that there are interactions between scheduled actions, the standard execution is stopped, until all interactions have been managed. On the other hand, the treatment of interactions between candidate actions is not necessary, since such actions will not necessarily have to be executed on the patient (their execution depends on the future status of the patient, after the execution of the scheduled actions). However, it is important that physicians are notified soon that such interactions may have to be faced in a near future.

3.5 Treatment of the Interactions

The management of interactions is performed by the physicians with the support of GLARE-SSCPM Interaction Management module. Given an interaction, such a model provides physicians with the possibility of choosing the most appropriate management, and helps them in its application to the original CIGs. Specifically, a result of the application of the *Interaction Management* module, the physician can see how the original CIGs are modified when applying the chosen interaction management operation to the given interaction. Such a process can be iterated, until one of the possible management is chosen by the physicians.

Notably, in case more than one interaction has to be managed, the Conciliation Module is invoked, in order to check the consistency of the different modifications to the original CIGs. In case they are consistent, the CIGs in CIG_{pat} are updated with the selected managements, and the Executor Modules can re-start execution on the updated CIGs. In case they are not consistent, physicians are requested to backtrack to the management of the interactions to consider alternative management options, until a consistent set of managements is determined.

3.6 Comorbidity Master Module

A new dedicated module has to be introduced, in order to support the run-time management of comorbidities. In the following, we informally describe it (called Comorbidity Master module).
The Comorbidity Master Module takes in input

(1) the clinical record of the patient
(2) the CIGs currently under execution (indicated by CIG_{pat} henceforth)
(3) the current action in each CIG in CIG_{pat}
and manages for each CIG_i in (2), two local data structures:
(4) the set SA_i of *scheduled* actions
(5) the set CA_i of *candidate* actions

When the Comorbidity Master Module receives in input (from the Executor of one of the CIGs) a message that the execution of an action A_h in the CIG CIG_k in CIG_{pat} has terminated, and the action A_h is not a decision action, it simply updates the set SA_k by deleting A_h from it.

On the other hand, in cases

(i) it receives in input a message that a new CIG CIG_k has been activated on the patient (so that CIG_k is added to CIG_{pat})

(ii) it receives in input a message that the execution of an action A_h in the CIG CIG_k has terminated, and A_h is a decision action, and $Path_j$ has been selected

several operations have to be performed, for the "run-time" identification and resolution of possible interactions. In such cases, the Comorbidity Master Module

1. evaluates the new set SA_i of scheduled actions (and their temporal windows)
2. evaluates the new set CA_i of scheduled actions (and their temporal windows).
3. Invokes the interaction detection module on the sets of scheduled actions of the CIGs in CIG_{pat} (not considering the decision actions). In case some interaction is detected, the set INT_sched of such interactions is sent to the Executors of the CIGs in CIG_{pat}, with the indication that such interactions occur between scheduled actions.
4. Invokes the interaction detection module on the sets of candidate actions of the CIGs in CIG_{pat} (not considering the decision actions). In case some interaction is detected, the set INT_cand of such interactions is sent to the Executors of the CIGs in CIG_{pat}, with the indication that such interactions may occur between candidate actions.

Notably, the detection of interaction is based on the Knowledge base, and is fully automatic.

In the following, we show two examples of Steps 1 and 2 above. Concrete examples of the management of CIG interactions have been reported in (Piovesan and Terenziani, 2015; Piovesan et al., 2018).

Example. Suppose that the CIG_1 is being executed on patient 1, and that, when A_2 is under execution (is current), the treatment of a new disease, through the CIG CIG_2, is started. The start of the execution of CIG_2 triggers the Comorbidity Master Manager for patient 1. $CIG_{pat1}=\{CIG_1,CIG_2\}$, and the set of scheduled and candidate actions are valuated as follows:

$SA_1=\{A_2,A_3,D_1\}$, $CA_1=\{A_4,T_2,A_9,T_3\}$

$SA_2=\{A_{100},A_{101},D_{100}\}$, $CA_1=\{T_{100},A_{104},A_{105},A_{106},D_{101},A_{107},A_{108}\}$

The Interaction Detection module is activated, and interactions between A_2,A_3,A_{100},A_{101} (if any) must be managed by physicians (while the interactions considering also A_4, A_9, A_{104}, A_{105},A_{106}, A_{107},A_{108} (if any) are pointed out to the physicians.

Example. Suppose that CIG_1 and CIG_2 are being executed on patient 1, that the current actions in CIG_1 and CIG_2 are A_3 and D_{100} respectively. We thus have

$SA_1=\{A_3,D_1\}$, $CA_1=\{A_4,T_2,A_9,T_3\}$

$SA_2=\{D_{100}\}$, $CA_2=\{A_{104},A_{105}, D_{101}, A_{106},A_{107}, A_{108}\}$

Suppose then that the execution of decision D_{100} give as result the path starting with A_{106}. Then,

$SA_2=\{A_{106},D_{101}\}$, $CA_2=\{A_{107}, A_{108}\}$

4 CONCLUSIONS

The CIG literature has devoted a considerable attention to the treatment of comorbid patients. However, the problem of supporting physicians in the "run-time" detection and management of CIG interactions has been quite neglected: in short, (Zamborlini et al., 2014) copes with knowledge-based interaction detection (but not with CIG merge), while the other approaches discussed in Section 1.2 focus on the merge of whole CIGs, assuming to have a pre-defined set of possible interactions, and of the way to treat each of them.

In this paper, we propose a comprehensive approach to run-time comorbidity management, based on GLARE and GLARE-SSCPM, which (i) automatically detects the "relevant" parts of the CIGs (i.e., scheduled and candidate actions), (ii) automatically detects possible interactions between them, (iii) supports physicians in the choice of the most appropriate management of such interactions. A prototypical implementation of the proposed approach is under development. Future works mainly concern a full realization of a tool, and an extensive experimentation on different concrete cases of comorbidity.

REFERENCES

Anselma, L., Terenziani, P., Montani, S., Bottrighi, A., 2006. Towards a comprehensive treatment of repetitions, periodicity and temporal constraints in clinical guidelines. Artif. Intell. Med., Temporal

Representation and Reasoning in Medicine 38, 171–195. https://doi.org/10.1016/j.artmed.2006.03.007

Bottrighi, A., Giordano, L., Molino, G., Montani, S., Terenziani, P., Torchio, M., 2010. Adopting model checking techniques for clinical guidelines verification. Artif. Intell. Med. 48, 1–19. https://doi.org/10.1016/j.artmed.2009.09.003

Institute of Medicine (US), 1990. Clinical practice guidelines directions for a new program. National Academy Press, Washington, D.C.

Montani, S., Terenziani, P., Bottrighi, A., 2005. Exploiting Decision Theory for Supporting Therapy Selection in Computerized Clinical Guidelines, in: Miksch, S., Hunter, J., Keravnou, E.T. (Eds.), Artificial Intelligence in Medicine, 10th Conference on Artificial Intelligence in Medicine, AIME 2005, Aberdeen, UK, July 23-27, 2005, Proceedings, Lecture Notes in Computer Science. Springer, pp. 136–140. https://doi.org/10.1007/11527770_19

Peleg, M., 2013. Computer-interpretable clinical guidelines: A methodological review. J. Biomed. Inform. 46, 744–763. https://doi.org/10.1016/j.jbi.2013.06.009

Piovesan, L., Anselma, L., Terenziani, P., 2015. Temporal detection of guideline interactions, in: HEALTHINF 2015. Presented at the International Conference on Health Informatics 2015, Scitepress, pp. 40–50. https://doi.org/10.5220/0005186300400050

Piovesan, L., Molino, G., Terenziani, P., 2014. Supporting Physicians in the Detection of the Interactions between Treatments of Co-Morbid Patients, in: Healthcare Informatics and Analytics: Emerging Issues and Trends. IGI Global, pp. 165–193.

Piovesan, L., Terenziani, P., 2016. A Constraint-Based Approach for the Conciliation of Clinical Guidelines, in: Advances in Artificial Intelligence - IBERAMIA 2016, LNCS. Presented at the Ibero-American Conference on Artificial Intelligence, Springer International Publishing, pp. 77–88. https://doi.org/10.1007/978-3-319-47955-2_7

Piovesan, L., Terenziani, P., 2015. A Mixed-Initiative approach to the conciliation of Clinical Guidelines for comorbid patients, in: KR4HC 2015, Lecture Notes in Artificial Intelligence. Springer International Publishing, Pavia, pp. 95–108.

Piovesan, L., Terenziani, P., Molino, G., 2018. GLARE-SSCPM: an Intelligent System to Support the Treatment of Comorbid Patients. IEEE Intell. Syst. https://doi.org/10.1109/MIS.2018.111144734

Riaño, D., Collado, A., 2013. Model-Based Combination of Treatments for the Management of Chronic Comorbid Patients, in: Artificial Intelligence in Medicine. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 11–16.

Sánchez-Garzón, I., Fernández-Olivares, J., Onaindía, E., Milla, G., Jordán, J., Castejón, P., 2013. A Multi-agent Planning Approach for the Generation of Personalized Treatment Plans of Comorbid Patients, in: AIME 2013, Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 23–27. https://doi.org/10.1007/978-3-642-38326-7_4

Sittig, D.F., Wright, A., Osheroff, J.A., Middleton, B., Teich, J.M., Ash, J.S., Campbell, E., Bates, D.W., 2008. Grand challenges in clinical decision support. J. Biomed. Inform. 41, 387–392. https://doi.org/10.1016/j.jbi.2007.09.003

Ten Teije, A., Miksch, S., Lucas, P. (Eds.), 2008. Computer-based medical guidelines and protocols: a primer and current trends, Studies in health technology and informatics. IOS Press, Amsterdam.

Terenziani, P., Molino, G., Torchio, M., 2001. A modular approach for representing and executing clinical guidelines. Artif. Intell. Med. 23, 249–276. https://doi.org/10.1016/S0933-3657(01)00087-2

Terenziani, P., Montani, S., Bottrighi, A., Torchio, M., Molino, G., Correndo, G., 2004. A context-adaptable approach to clinical guidelines. Stud. Health Technol. Inform. 107, 169–173.

Wilk, S., Michalowski, W., Michalowski, M., Farion, K., Hing, M.M., Mohapatra, S., 2013. Mitigation of adverse interactions in pairs of clinical practice guidelines using constraint logic programming. J. Biomed. Inform. 46, 341–353. https://doi.org/10.1016/j.jbi.2013.01.002

Zamborlini, V., da Silveira, M., Pruski, C., ten Teije, A., van Harmelen, F., 2014. Towards a Conceptual Model for Enhancing Reasoning About Clinical Guidelines, in: Miksch, S., Riaño, D., ten Teije, A. (Eds.), Knowledge Representation for Health Care. Springer International Publishing, Cham, pp. 29–44.