2

TMR: A CORE CONCEPTUAL MODEL FOR CLINICAL GUIDELINES
- *APPLIED TO COMORBIDITY ANALYSIS*

No problem can be solved from the same level of consciousness that created it.

Albert Einstein

Computer-Interpretable Guidelines (CIGs) are representations of Clinical Guidelines (CGs) in computer interpretable languages. CIGs have been pointed as an alternative to deal with the various limitations of paper-based CGs on supporting healthcare activities. Although the improvements offered by existing CIG languages, the complexity of the medical domain requires advanced features in order to reuse, share, update, combine or personalize their contents. We propose a conceptual model for representing the content of CGs as a result from an iterative approach that takes into account the content of real CGs, CIG languages and foundational ontologies in order to enhance the reasoning capabilities required to support several CG-tasks. In particular, we apply our approach to the task of comorbidity analysis, illustrate the model with a realistic case study (Duodenal Ulcer and Transient Ischemic Attack) and compare the results against an existing approach.

This chapter is based on *Zamborlini, V.; da Silveira, M.; Pruski, C.; ten Teije, A.; van Harmelen, F. "Towards a Conceptual Model for Enhancing Reasoning about Clinical Guidelines: A case-study on Comorbidity", in Knowledge Representation for Health-Care (KR4HC). Lecture Notes in Computer Science, vol. 8903, Vienna, Austria: Springer International Publishing, 2014, pp. 29-44.*

2.1 INTRODUCTION

Clinical guidelines (CGs) assemble statements provided by the best available evidences. Their goal is to assist healthcare professionals on the definition of the appropriate treatment and care for people with specific diseases and conditions. Formal representations for CGs, called computer-interpretable guideline (CIGs), have been proposed to overcome some limitations of paper-based CGs using dedicated languages (e.g., PROforma [85], GLIF [16], Asbru [60]). They can be integrated to health information systems to support healthcare professionals in their daily practice. Although being expressive, existing CIG specification languages are designed for one main objective: to execute the CIG into a treatment or diagnosis plan.

However, the evolving requirements from the medical field combined with the properties of information systems, demand other advanced features. New requirements are motivated to tackle problems like comorbidity (combining guidelines to define appropriate treatments for patients suffering from several diseases), CG update (taking into account new findings from clinical studies) or treatment personalization (taking into account patient's preferences).

To cope with these kind of problems, CIGs must be improved in order to offer more *reasoning capabilities*. For instance, consider a patient that suffers from Duodenum Ulcer (DU) and from Transient Ischemic Attack (TIA). Two different guidelines need to be combined to define a treatment. But, a closer analysis of them shows that these guidelines lead to adverse interactions when combined. Combining CIGs, detecting conflicts, and including new information have not been the focus of existing CIG languages and their underling editing and execution tools. Therefore, a representation language is needed that enables reasoning over CG information for several tasks like combining or updating CIGs.

In this paper, we introduce a new conceptual model to enhance the reasoning capabilities of CIGs. The elements of the proposed model are identified following an iterative approach to explicitly represent the semantics of recommendations and medical actions. The reasoning capabilities of the proposed model have been assessed on a realistic case study dealing with conflicts detection and solving in case of comorbidity. The remainder of the paper is structured as follows: Sect. 2.2 presents the analysis of the related work. In Sect. 2.3 we propose a conceptualization of our model before applying it to the comorbidity use case in general, and then to a particular case study (stroke + ulcer). In Sect. 2.4 we discuss the results and future work and wrap up with concluding remarks in Sect. 2.5.

2.2 RELATED WORK

Several CIG description languages are proposed in the literature. They provide different methods to model the content of CGs into CIGs. Studies comparing these languages highlighted the qualities and the scope of each one [46, 66]. They mainly analysed three aspects: (1) the edition and execution of CIGs, (2) the capacity to collaborate with other systems, and (3) the dissemination properties. Isern & Moreno [46] centred their study on the editing and execution tools. They underline that the interoperability between systems is the most important barrier to overcome in order to promote CIGs. A standard description language and a standard electronic health record (EHR) would help the progress in this domain and avoid development of ad hoc solutions.

However, Peleg [66] pointed out the difficulty to define a standard language that integrate the different components of each language, and proposes to start by splitting CIGs into small size "knowledge chunks". She argues that defining small chunks of decision logics will contribute to cope with three complex and important problems: sharing/reusing, combining and maintaining knowledge. In this paper, we propose a model that is meant to address those problems, though we focus on the comorbidity issue.

With the increasing of aged population and the frequency of comorbidities, this subject has been considered as an important topic of research in the medical domain. Consequently, there is a high demand for computer systems that support medical researches in comorbidity. Recent publications propose semi-automatic combinations of CIGs, some of which we summarize hereafter. Authors claim that existing languages were not designed to address this problem and they propose new CIG representation formalisms for it.

Jafarpour & Abidi [49] adopted OWL to describe CIGs. They also built a merging representation ontology to capture merging criteria in order to achieve the combination of CIGs. SWRL rules were used to identify potential conflicts during the merging process. All conditions related to the merging process need to be described by the rules, increasing the effort to maintain the system up-to-date, and reducing the possibility of sharing knowledge. However, some related problems were not yet (completely) addressed in their work, for instance, potential contradictions between rules, the scalability of the merging model to combine several CIGs, and how the ontology/rules are maintained up-to-date.

A different approach was proposed by Wilk et al. [106]. They describe CIGs as an activity graph and propose to use constraint logic programming (CLP) to identify conflicts associated with potentially contradictory and adverse activities resulting from applying two CGs to the same patient. The goal is to use this approach to alert physicians about potential conflicts during the definition of the treatment plan. The temporal aspect is not considered, thus the approach can only be applied to specific situations (e.g. acute diseases diagnosed during a single patient-physician encounter). Although their model allows reasoning over a subset of the CIGs content (the conditions) and propose possible conflict solutions, the whole work of combining CIGs remains manual. This approach also considers that all predicates use the same terminology and that they can have only two states (true or false). The case study used to demonstrate the applicability of the approach in [106] shows the complexity of combining CIGs and the necessity of external knowledge sources for taking decisions. Inspired on this case study we evaluate the applicability of our model in the comorbidity analysis task.

Another method to address the CIGs combination problem is proposed by Riano & Collado [75]. They define a language to describe CIGs as actions blocks and decision tables. A generic treatment model is proposed to decide which action is appropriate to a chronically comorbid patient, taking into account three criteria: seriousness, evolution, and acuteness. The expressivity of this language is intentionally limited in order to have a lightweight decision system. The combination of CIGs is the result of pairwise combination of CIGs entities (i.e., actions and decisions table) according to a set of rules that allow identifying conflicts and reorganising or merging actions (in specific and predefined situations). The simplified CIGs representation and the specification of more general rules (for merging tasks) increase the reasoning capability of the system and reduce the maintenance work effort. However, reorganising care actions can raise some problems, especially those related to the clinical validity of modifications. In this case, the evidence-based medicine must be assured in the rules of the generic treatment model. An alternative to this problem is to associate intentions and goals to the actions, as proposed by Latoszek-Berendsen et al. [55]. However, they do not consider combining CIGs and evaluating the role of intentions in this process.

The idea of evaluating pairwise actions associated to goals is exploited in the work of Sanchez-Garzon et al. [79]. They adopt the HTN plan description language to describe CIGs, and they use multi-agents techniques to generate treatment plans and identify potential conflicts between care actions. Treatment goals are considered to solve conflicts, but the assumption of all effects of an action is observed in the patient (and included in the patient data) limits the applicability of their approach. A probabilistic representation of effects would be closer to observations from evidence-based studies, but it would increase the complexity of the reasoning. Although the good preliminary results claimed by the authors, the low interoperability and the complexity of maintenance of agents has been underlined in several publications as a challenge of the domain.

In the referred approaches the care actions are represented as textual information (or labels) and their semantics is not clearly defined, for example, "*Start Aspirin*" and "*Stop Aspirin*" are represented as unrelated actions, what confirms the outcomes of Bonacin et al. [13]. Consequently a specific rule is required to define them as conflicting actions, while it could be automatically detected by reasoning over the meaning of the actions.

Moreover, few evidences about how these actions impact the patients' health state are formalized. For instance, the intention of an action for a specific treatment, their potential effects (desired and side-effects) and the situation (describing the context). Understanding the semantics of the care actions and the related impacts is considered as an important source of information to increase the reasoning capabilities and better explain the causes of conflict [13].

Another potential advantage of having less constraints and more detailed actions is the reduction of required maintenance efforts. New findings about one action can easily be integrated to the CIGs without requiring a whole analysis of the impact of these changes. Collaborative work to specify care actions can also promote the reuse of knowledge chunks, facilitating CIGs construction/update. In this paper we aim to provide a more detailed semantics for care actions and recommendations, and to evaluate the benefits for the task of comorbidity analysis.

2.3 THE TMR MODEL

We present in this section the Transition-based Medical Recommendation (TMR) Model for Clinical Guidelines, a conceptual model designed to capture the core knowledge structure for CGs' recommendations. The purpose is to favor the reasoning capabilities required by different CIG tasks, like combining CIGs to deal with comorbidity. On what follows we present the conceptualization adopted for our model and its application to the comorbidity analysis.

2.3.1 *Conceptualization*

In order to investigate the knowledge structure in the CGs domain, we adopted an approach that involves studying several CGs, CIG languages, CIG use-cases and foundational ontologies. We adapted two example recommendations from a CG for Peptic Ulcer[1] to illustrate the concepts and issues to be handled:

Section V

1. For patients with ulcer not associated with Helicobacter Pylori (HP), maximal dose of proton pump inhibitors (PPI) is recommended;
2. For patients with ulcer caused by NSAID (non-steroid anti-inflammatory drugs), NSAID use should be discontinued.

According to Peleg and colleagues [66], all current GIG languages provide some structure for representing **Actions** and **Decisions**. Considering a structure "if ... then ..." for representing the decision and the corresponding action, a representation for the mentioned example would be: (1) *if "ulcer is not caused by HP" then "administer PPI on maximum dose"*; and (2) *if "ulcer is caused by NSAID" then "do not administer NSAID"*. While the **Actions** represent the tasks described in a CG, the **Decisions** regard mainly the evaluation of context (Pre-Situations) that would enable to choose the appropriate actions. Moreover, few languages also provide support for expressing the potential effects of actions (Post-Situations) like Asbru and Proforma.

Some representation issues can be observed in the aforementioned example: (i) how to identify and represent the information that is

1 http://www.aiha.com/en/WhatWeDo/PracticeGuidelines_CPGPI.asp

implicit in the CG text itself, like the expected outcome for a recommended action; and (ii) how to represent "negative" actions such as in the example recommendation 2. A proper solution for these issues may enhance the capability of reasoning over the knowledge structure (the dosage is out of the scope in this work).

In order to guide our interpretation of the CG knowledge structure we use foundational (top-level) ontologies (such as UFO [29]) that define generic entities and its relations, e.g. actions and situations. Those theories provide means to justify the modeling choices made in a model. Although the study of those theories is an important part of our approach, it is not the goal of this paper to provide a precise ontologically-founded definition for the concepts.

In this work we select some entities in CG context as a small/-core knowledge chunk to be analyzed and combined to represent more complex scenarios. The main concepts adopted in the TMR model for CG domain are summarized in Table 2.1, namely Situation Type, Care Action Type, Transition and Recommendation. We consider those concepts as being atomic, since the study of their compositionality is not in the scope of this work.

The aforementioned example is instantiated in Fig. 2.1 according to the TMR model, also considering relevant implicit information.

Table 2.1: TMR Concepts Summary

Situation Type	Represents a property, which characterizes a patient, and its admissible values
Care Action Type	Represents the action types that can be performed by healthcare agents in order to change a situation.
Transition	Represents the possibility of changing a situation regarding a patient by performing a care action type.
Recommendation	Represents a suggestion to either pursue or avoid a transition promoted by a care action type.

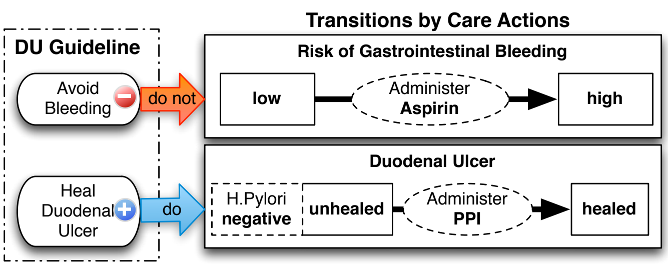

Figure 2.1: Instance schema for the TMR model

The bigger rectangles (boxes) represent the transitions regarding a the possibility of a situation to be changed by executing a care action (e.g. *risk of gastrointestinal bleeding* changes from *low* to *high* by *administering aspirin*). The care action is represented as an ellipse within that box (e.g. *administering aspirin*). The pre- and post-situations of the transitions are expressed in smaller boxes connected by a directed arrow, which goes through the care action (e.g. *risk of gastrointestinal bleeding low* or *high*). A situation type that is not changed by the transition is expressed within a dotted box before the box representing the transformable situation (e.g. *H.Pylori negative*). Finally, a recommendation is represented as a rounded box, with a label inside for reference (e.g. *heal duodenal ulcer*). It is connected to a transition by a thick labeled arrow indicating either the transition (i) to be pursued with a label *'do'* (also highlighted by blue color and a '+' sign) or (ii) to be avoided with a label *'do not'* (also highlighted by red color and a '-' sign). Therefore, the recommendation named '*heal duodenal ulcer*' recommends for *patients with H.Pylori negative* to pursue the transition promoted by the action '*administer PPI*' in order to have the *duodenal ulcer* transformed from *unhealed* to *healed*, while the recommendation named '*avoid bleeding*' recommends to avoid the transition promoted by the action '*administer aspirin*' in order to not have the *risk of gastrointestinal bleeding* transformed from *low* to *high*.

We hereafter explain our modeling choices. Firstly we distinguish between instance (individual) and type (universal) levels. The instance level regards, for example, the action occurrence '*John takes PPI*' that leads from a pre-situation '*John with ulcer*' to a post-situation '*John without ulcer*'. The recommendations in CGs, however, do not regard the instance level, i.e. the factual situations and action occurrences, but the **Type Level**, i.e. the **Care Action Types** and **Situation Types**, as well as the relations between them. An example of care action type is '*administer aspirin*', which can be performed by healthcare agents such as a physician, a nurse, or the patient itself, while an example of situation type is '*patient with risk of gastro-intestinal bleeding (GIB)*'.

If in the one hand an action occurrence directly relates pre- and post-situations according to the promoted change, on the other hand an **Action Type** is expected to be related with one or more pairs of **Pre-/Post-Situation Types**. Indeed, Textor [87] mention the need for a space of outcomes for an action type (e.g. throwing a dice have 6 possible outcomes). Although in the medical domain the outcomes of an action type usually cannot be precisely and completely defined, they constitute the core knowledge that underlies the clinical recommendations. Indeed, the different changes that can be promoted by a care action type must be taken into account as desired or side-effects for a patient (type). For example, '*administer aspirin*' has two possible effects: *anti-prostaglandin* (anti-inflammation, fever-reducing, pain reliever) and *anti-platelet* ('blood thinner') agent. By inhibiting the formation of prostaglandins, aspirin deplete the protective barrier in the stomach against the acid substances, leading to peptic ulcers. Thus, for patients with bleeding risks or duodenal ulcer, aspirin may have a negative effect, while for patients with cardiovascular events risk, aspirin may have a positive impact.

Aligned to this idea, we introduce the concept **Transition** to relate a care action type to pre-/post-situation types and represents the possibility of achieving that change by performing the referred action. Thus, by assigning different transitions to a care action type, we define its '*space of transitions*'. Finally, the **Recommendation** can be seen as a commitment for healthcare agents to either pursue or avoid a transition, whilst the **Guideline** contains a set of recommendation about transitions.

Moreover, the situation types involved in a transition can be classified as: (i) **Non-Transformable Situation Type** regards a property that is not to be changed in a transition, but is needed as a filter condition (*patient is a woman*); (ii) **Transformable Pre-Situation Type** regards a property and value that is to be changed in a transition (*ulcer is unhealed*); (iii) **Post-Situation Type** regards the expected value for the property that is to be changed in a transition (*ulcer is healed*).

The aforementioned concepts and relations are represented in an UML class diagram in Fig. 2.2. While one guideline is composed of one or more recommendations, the recommendation is part of one guideline. A recommendation either recommends to pursue or to avoid one transition. The latter is promoted by one care action type, which in turn can promote one or more transitions. Situation types can be pre- or post-situation type in the context of different transitions, which must have one transformable situation type, one expected post-situation type and may have as filter condition some non-transformable situation types.

Finally, the situation types can also be classified either from the perspective of the patient health condition or of the HealthCare System (HCS) as follows: (i) **Patient Health Condition Type**: regards

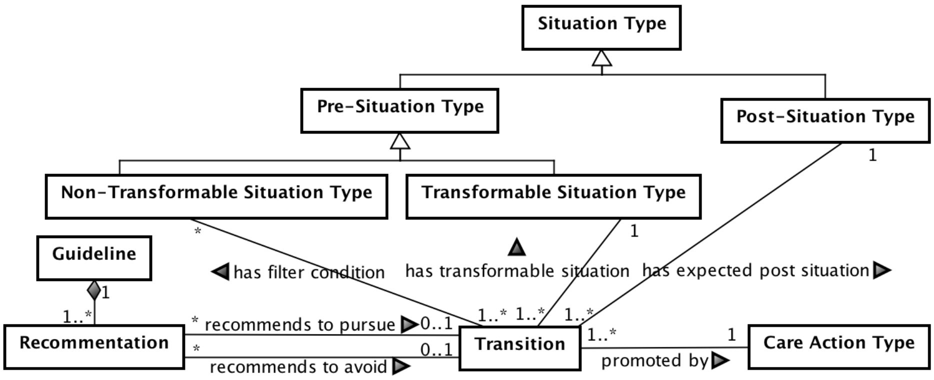

Figure 2.2: UML class diagram for the TMR Model

the properties that define the patient health condition (e.g. *patient has ulcer*); (ii) **HCS Epistemic State Type**: regards the knowledge about the patient properties by the HCS (e.g. *H. Pylori presence is unknown*); and (iii) **HCS Patient Status Type**: regards the status of a patient in a HCS (e.g. *patient is forwarded*). The transitions regarding these situation types can be classified according to the same criteria, as well as the action type that promotes the transition and the recommendation itself. The concepts here defined are further illustrated in the case study presented in Sect. 2.3.3.

2.3.2 *TMR Application to Comorbidity Analysis*

We evaluate the proposed model by analyzing CIGs combined due to comorbidity, which regards taking into account two diseases that patients may suffer from simultaneously. If this issue is not correctly addressed a patient will possibly have an inadequate treatment. As a consequence it is necessary to combine and analyze CIGs and/or treatment plans related to the different diseases in order to identify and solve the issues that may appear in the process of treating comorbid patients.

As mentioned in Sect. 2.2, since the current CIG languages do not properly address the comorbidity analysis, some approaches have being proposed to this end. Jafarpour & Abidi [49] mention two classifications for the existing approaches, namely: (i) Pre-Execution Level Merging: issues are handled during the treatment prescription; and (ii) Execution Level Merging: issues are handled after the treatment prescription. We introduce here an extension for this classification as follows:

GUIDELINE-LEVEL VERIFICATION aims to handle the combining issues at the guideline level (before execution). The result is a combined version of CIGs in which guideline-level issues are addressed. (e.g. in [106] the authors combine the CIGs before executing, though their goal is to produce a treatment for a specific patient).

ON-PRESCRIPTION VERIFICATION aims to handle the combining issues during the prescription of the treatment. The result is a merged treatment free of treatment-level issues. It can be applied between CIGs or between CIGs and existent treatments (e.g. [79]).

AFTER-PRESCRIPTION VERIFICATION aims to handle the combining issues among treatments. The result is a merged treatment applicable free of treatment-level issues (e.g. [75]).

ON-TREATMENT-EXECUTION VERIFICATION aims to handle the issues that cannot be foreseen, since they happen during the treatment execution. The result can be an alert to interrupt the treatment execution (e.g. [14]).

We believe that these types of approaches are complementary, since on the one hand it is useful to anticipate the issues when possible, but on the other hand it is complex (maybe not possible) to anticipate all of them. The work presented in this paper fits to the Guideline-level Verification, since we aim to produce a combined version of CIGs that addresses guideline-level issues and can be applied to many patients.

A simple scenario for *comorbidity analysis* is presented in Figure 2.3 according to the TMR model. When the recommendations from

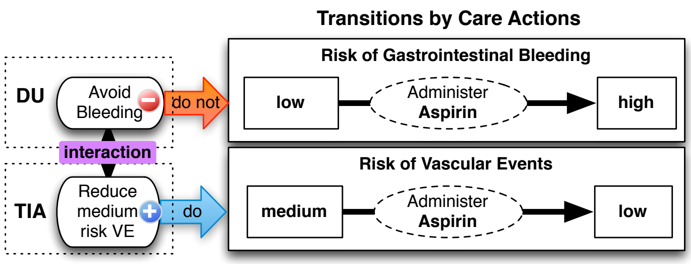

Figure 2.3: Comorbidity example according to the TMR model

DU CIG and TIA CIG are combined, it can be identified an interaction between the recommendations '*avoid bleeding*' and '*reduce medium risk VE (vascular event)*'. The interaction is represented as a label connected via a small arrow to the recommendations. In this case the interaction regards recommending to pursue and to avoid transitions promoted by the same care action type, namely, '*administer aspirin*'.

Note that by applying the TMR model it is possible to detect interactions among recommendations, but not yet the conflicts. In order to identify conflicts, we would need both: (i) check if the interaction is unavoidable, i.e. no alternative path that can be derived (for the same purpose/context) and (ii) consult external knowledge base in order to check for overdoses or incompatibilities. However, the scope of this paper is restricted to identify the interactions, which could lead to conflicts or require attention from the experts. Moreover, we consider that the interactions are not all unwelcome (e.g. the recommendations to inverse transitions may be desirable and the alternative ones are useful to avoid conflicts) although they could still require some attention (e.g. defining which alternative recommendation is preferred). Therefore, we distinguish the following types of interaction among recommendations:

CONTRADICTION: when a set of recommendations that can lead to an undesired (non-recommended) final situation.

Opposed recommendations to the same care action: when transitions promoted by a same care action are recommended to be pursued in a CG and avoided in another, i.e. the execution of a care action may lead both to a desired and an undesired post-situations (e.g. administering aspirin reduce the risk of vascular events but also increase the risk of gastrointestinal bleeding).

Opposed recommendations to similar transitions: when a situation is the post-condition of transitions promoted by different care actions that are recommended to pursued and to be avoided, i.e. the execution of a care action will promote a post-situation that had also been stated as undesired (e.g. recommending '*administer ACE inhibitor*' to pursue *lower blood pressure* while also recommending to avoid similar effect promoted by '*administer beta-blockers*').

Recommendations to inverse transitions: two recommendations that revert each other effect (e.g.: '*administer midodrine*' is recommended to *increase blood pressure* and *''administer ACE inhibitor'* to *decrease it*).

REPETITION: when a set of recommendations can lead to overdose, and are susceptible to optimization.

Repeated recommendations to the same care action: More than one recommendation to pursue transitions promoted by one care action (e.g.: *'perform blood exam'* is recommended twice).

ALTERNATIVES: when a set of recommendations holds as alternatives.

Recommendations to pursue similar transitions promoted by different care action: recommendations that can promote similar effects (e.g. both *'administer aspirin'* and *'administer clopidogrel'* may *reduce the risk of vascular events*).

REPAIRABLE: when a set of recommendations has one promoted (undesired) effect that can be undone by another recommendation.

Transition recommended to be avoided whose inverse transition is recommended to be pursued: when the undesired effect of a transition can be undone by another recommended transition (e.g. the undesired effect *'increase the risk of gastrointestinal bleeding'* promoted by *'administer aspirin'* can be undone by *'administer PPI'*, which *decreases that risk*).

We compared the aforementioned classifications with the ones proposed in GLINDA Project[2]. For example, the *Opposed recommendations to similar transitions* could be mapped both to GLINDA *Cumulative Number Constraint* and *Inconsistent Goals*. We intend to further investigate the matching to the GLINDA classification for conflicts.

2.3.3 *Evaluation on Comorbidity Case Study*

In this section we apply our model to a case study on the comorbidity analysis task. We repeat the experiment done by Wilk et al. [106] by modeling the CGs for Duodenal Ulcer (DU) and Transient Ischemic Attack (TIA) and merging them into a combined DU-TIA CIG. However, since the CIGs presented in the referred work do not provide all

2 http://glinda-project.stanford.edu/guidelineinteractionontology.html

information that we need in the TMR model, we made some assumptions based on related CGs or common sense. Figure 2.4 presents the DU CIG represented according to both [106] and the TMR model.

The action '*Stop aspirin if used*' in the original CIG is represented in the TMR CIG as a recommendation named '*Avoid Bleeding*' that recommends to avoid a transition promoted by the care action '*Administer Aspirin*'. The undesired transition can lead from the situation '*low risk of gastrointestinal bleeding*' to '*high risk of gastrointestinal bleeding*'. The decision point '*H.Pylori test?*' in the original CIG is separated in the TMR CIG as: (i) a (epistemic) recommendation to the transition promoted by the care action '*H.Pylori Exam*' that reveals the result for the H.Pylori infection (from *unknown* to *known*); and (ii) filter pre-situation types that would enable one of two recommendations named '*healing DU*'. When '*H.Pylori is positive*' the care action '*Eradication Therapy*' can lead from the pre-situation '*DU is unhealed*' to the post-situation '*DU is healed*'. When '*H.Pylori is negative*' instead, the care action '*Administer PPI*' can promote the same effect. The two recommendations aforementioned represent the actions '*Start Eradication Therapy*' and '*Start PPI*' from the original CIG. A similar procedure were applied for the other actions and decision points.

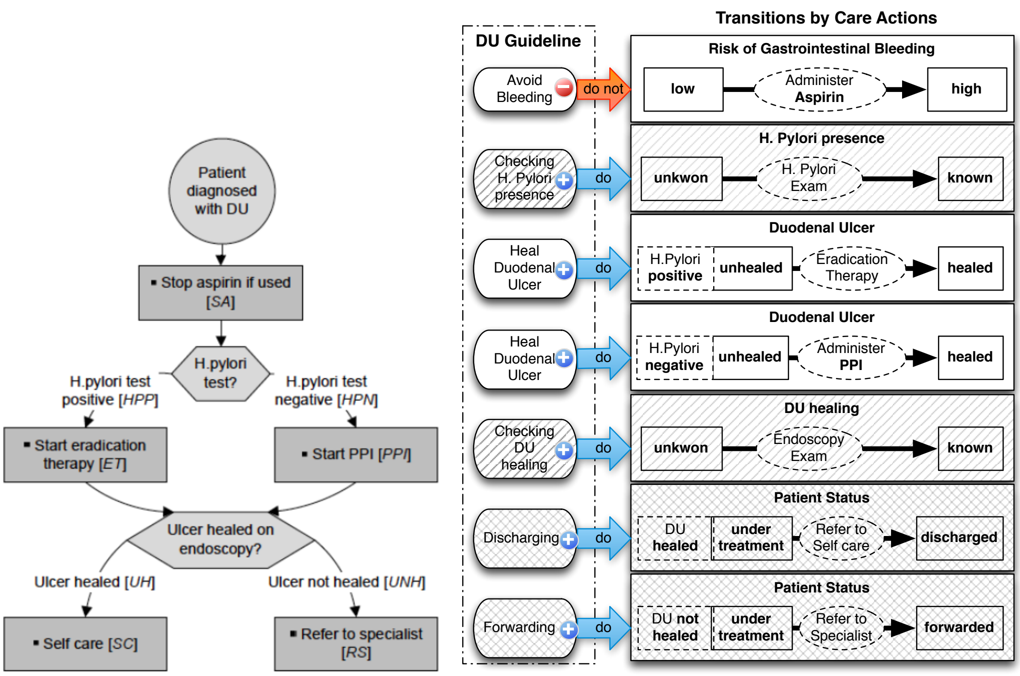

Figure 2.4: DU CIG according to [106] (left side) and to TMR Model (right side)

In addition, different classifications for the Situations Types are distinguished in Fig. 2.4 (right side) by different backgrounds: (i) Patient Health Conditions - filled background; (ii) HCS Epistemic Situations - diagonal lines background; and (iii) HCS Patient Status - crossing-lines background. Corresponding classifications for transitions, actions and recommendations follow the same pattern in the figure.

Figure 2.5 presents the TIA CIG represented according to both [106] and the TMR model. We represent for the TMR model only two recommendations regarding Health Condition Transitions that are relevant for this case study (highlighted in the left side of Fig. 2.5). The actions '*Start Aspirin*' and '*Start Dipyridamole*' in the original CIG are represented as the recommendations named '*Reduce Medium Risk VE*' and "*Reduce High Risk VE*". They recommend respectively the transitions promoted by the care actions '*Administer Aspirin*' that changes the '*risk of vascular events*' from *medium* to *low*, and the transition promoted by the care action '*Administer Dipyridamole*' that changes the '*risk of vascular events*' from *high* to *low*.

Finally, when combining the two CIGs, the authors in [106] identified a conflict by consulting a restriction expressed in a Medical Background Knowledge (MBK). It states that the actions '*Stop aspirin if used*' and '*Start Aspirin*' cannot coexist, which indeed occurs in the combined version CIGs. In order solve the conflict, the authors had two possibilities also derived from the MBK: (i) to substitute aspirin

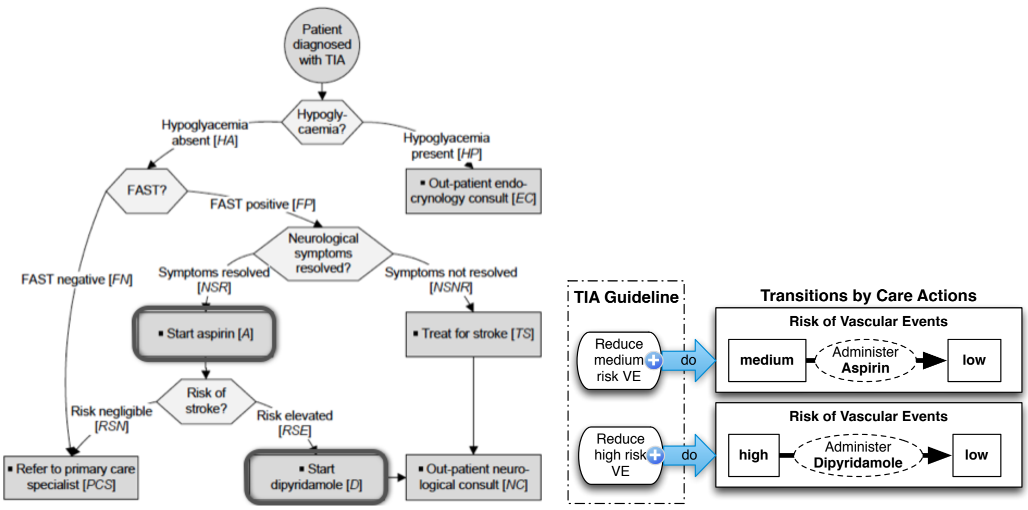

Figure 2.5: TIA CIG according to [106] (left side) and to TMR Model (right side)

by clopidogrel; and (ii) to combine aspirin treatment with PPI. They choose the second option and therefore introduced it in the merged CIG as '*Start PPI*' when the risk of stroke is elevated. They also excluded the recommendation '*Stop aspirin if used*' in order to avoid the aforementioned conflict. Since their final goal was not to produce a generic combined version of guidelines, but to prescribe a treatment for a specific patient, they proposed a solution that is applicable to a specific patient.

In its turn, the TMR model allows for identifying interactions among recommendations, depicted in Fig. 2.6. Firstly, the **contradiction** interaction between the recommendations '*Avoid bleeding*' and '*Reducing medium risk VE*' is identified, since they regard recommending to pursue and to avoid transitions that are promoted by the same action type (administer aspirin), highlighted in Fig. 2.6a. Then both mitigation alternatives proposed in [106] are introduced in the combined CIG, without excluding the 'conflicting' recommendation '*Avoid Bleeding*'. The alternative recommendations are named '*Protecting Duodenum*' and '*Reduce medium risk VE*' as highlighted in Fig. 2.6b. Finally, new interactions are identified between the original recommendations and the ones introduced as mitigation alternatives. They are depicted in Fig. 2.6b as: **alternative** interaction between the two recommendations named '*Reduce medium risk VE*' by administering either aspirin or clopidogrel; **repairable** interaction between the recom-

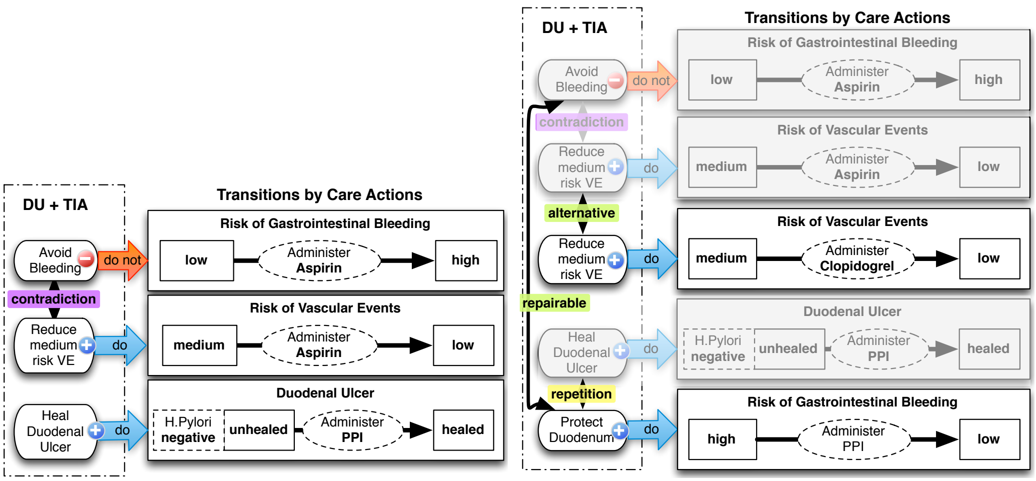

Figure 2.6: The left side (a) presents a (partial) combined DU+TIA CIG according to TMR model where a contradiction interaction is highlighted. In the right side (b) alternative recommendations are introduced and new detectable interactions are highlighted.

mendations '*Avoid Bleeding*' and '*Protect Duodenum*', as the latter can undo the effect of the former; and (iii) **repetition** interaction between the recommendations '*Protecting Duodenum*' and '*Heal Duodenal Ulcer*', since they are both recommend the action '*Administer PPI*'.

Therefore, the combined DU-TIA CIG produced using the TMR model does not eliminate the original conflict but allow to introduce both mitigation alternatives as recommendations. Actually the recommendation '*Avoiding bleeding*' about '*Administer Aspirin*' is not eliminated since it is a restriction that holds for DU patients regardless to what else disease they could have. Indeed, the resultant CIG is designed with the purpose of both (i) being applicable to many patients and (ii) being liable to further combination with other guidelines or treatments that a patient already follows. Finally, the interactions can be identified by relying on the described semantics for the referred care actions without consulting a MKB.

2.4 DISCUSSION AND FUTURE WORK

In this paper we propose the TMR model with the purpose of addressing other CIG tasks rather than CIG execution. The model is applied to the comorbidity analysis task and the results are compared to an existing approach presented in [106]. On what follows we discuss the proposed model, its positive aspects, limitations and future issues to be addressed according to the following perspectives: (i) the model itself (Sect. 2.4.1) and (ii) its application to the comorbidity analysis task (Sect. 2.4.2).

2.4.1 *The TMR Model*

The main contribution of the TMR model consists in a core knowledge structure for CGs that explicitly represents both (i) the care action types with the possible transitions between situations types that can be promoted and (ii) the recommendations as declarative suggestions to pursue or avoid such transitions. We advocate that the TMR model, by providing a detailed semantics for a small "CG knowledge chunk", can be a step towards addressing important issues like sharing/reusing, combining and maintaining knowledge such as argued by Peleg [66]. Although there is still place for investigation, we achieved improvements on addressing comorbidity analysis at the guideline level (discussed in Sect.2.4.2). We intend to apply the TMR model to other tasks such as adapting and updating CIGs and to analyze through the results the applicability of the current model and required adaptations.

Unlike in most CIG languages, the TMR model does not define a sequence among the recommendations, but further investigation on this issue is necessary. Indeed, while for some recommendations sequence is not necessary or desirable (e.g *do not administer aspirin*), for other ones the sequence can be derived by matching Post- and Pre-Situation Types (e.g. *If H.Pylori is negative then administer PPI for healing the DU* and *If DU is healed then discharge the patient*). We also reconsider other two common constructs of current CIG languages, namely the Decision Point and Enquiry (demand of information). The former is implicit in the evaluation of the pre-situations, while the enquiry is represented as a recommendation regarding the HCS Epistemic State. We intend to investigate how to address the known/unknown values for epistemic situations.

We intend to pursue compatibility with current CIG approaches by studying their underlying models and checking for a possible mapping to the TMR Model. In particular, the SDA approach by Riano [74] proposes a non-deterministic model for CIG that is composed of States, Decisions and Actions (SDA), but which is not meant for representing the semantics of the actions. We also plan to use biomedical terminologies/ontologies (e.g. SNOMED, ICD) for the descriptions Care Action and Situation Types.

Further improvements that we intend to investigate are (i) hierarchy and compositionality of the situations, actions, transitions and recommendations, (ii) the inclusion of new concepts (specially goals), (iii) the study of the recommendations as commitments and (iv) addressing temporality, probability and other features that characterize the domain and can enrich the TMR model. In summary, our future work will iteratively (re)apply improved versions of the TMR model to CG-tasks. Besides extending the TMR model, two important goals are: (1) providing formalized version of the TMR model and the reasoning concerning comorbidity analysis and (2) an implementation of them.

2.4.2 *Application to Comorbidity Analysis*

We applied the TMR model to the comorbidity analysis task and evaluated it by comparing with a related work [106]. We classify our approach as begin designed to address the combining issues at the guideline level, i.e. to produce a combined version of the CIG that can be applied to many patients and further combined with other CIGs. Then we explore the ability to identify interactions among recommendations, which could lead to conflicts or require attention from the experts, by relying on the CIG internal information rather than depending on external knowledge bases. Table 2.2 summarizes the comparison with the related work considering different aspects.

As future work on comorbidity analysis we intend to (i) investigate the formalization/automatization for identification of interactions and conflicts, as well as suggesting solutions, (ii) reapply improved versions of the TMR Model (according to the previously improvements mentioned) and (iii) evaluate it on more case studies. In particular, we intend to investigate and evaluate the ability to identify interactions among several recommendations in several CGs in the context of multimorbidity (interaciton analysis for more than two diseases).

2.5 CONCLUSION

The main contribution of this paper is the TMR model for representing CGs. This core model enhance some reasoning capabilities with respect to the current CIG languages, which are important to address the comorbidity analysis and possibly other tasks rather than CIG execution. It explicitly represents both (i) the (space of possible) transi-

Table 2.2: Comparison to a related work

	Wilk et al. [106]	**TMR Model**
Core Concepts	Actions & Decisions	Actions, Situations, Transitions, Recommendations
Description of Care Actions	Abstract/textual, does not favor reasoning	Detailed, favor reasoning
Knowledge	Procedural	Declarative
Format	Sequenced Actions & Decisions	Non-Sequenced Recommendations
Language	Workflow & CLP	Graphical notation
Combining Issues	Use an MKB for identifying and solving conflicts	Interactions among recommendations can be identified without MKB
Purpose	Introduce ONE alternative to produce a combined CIG for a SPECIFIC patient	Introduce MANY alternatives to produce a combined CIG applicable for MANY patients

tions between situations types promoted by the care actions types and (ii) recommendations as declarative suggestions to pursue or avoid transitions. By reasoning over such knowledge structure we are able to demonstrate improvements on addressing the comorbidity analysis. We evaluated the approach by repeating an experiment from the literature and comparing the results. We intend to iteratively improve the model and evaluate it by (re)applying it to other case-studies as well as other CG tasks (such as sharing, reusing, adapting and updating) at both conceptual and formal levels.