



MitPlan 2.0: Enhanced Support for Multi-morbid Patient Management Using Planning

Martin Michalowski¹(✉), Malvika Rao², Szymon Wilk³, Wojtek Michalowski²,
and Marc Carrier⁴

¹ University of Minnesota, Minneapolis, MN 55455, USA
martinm@umn.edu

² University of Ottawa, Ottawa, Canada

³ Poznan University of Technology, Poznan, Poland

⁴ The Ottawa Hospital Research Institute, Ottawa, Canada

Abstract. The complexity of patient care is growing due to an ageing population. As chronic illnesses become more common, the incidence of multi-morbidity increases. Generating disease management plans for multi-morbid patients requires the integration of multiple evidence-based interventions, represented as clinical practice guidelines (CPGs), that are designed to treat a single condition. Our previous work developed a mitigation framework called MitPlan that represented the generation of treatment as a planning problem. The framework used the Planning Domain Definition Language (PDDL) to represent clinical and patient information needed to identify and mitigate adverse interactions resulting from the concurrent application of multiple CPGs for a given patient encounter. In this paper we describe MitPlan 2.0 that supports shared decision-making by identifying a treatment plan optimized according to patient preferences, treatment cost, or perceived patient's adherence to medication. It mitigates adverse interactions using planning constructs, eliminating the need for procedural handling of adverse interactions and as such provides flexible and comprehensive decision support at the point of care. We demonstrate MitPlan 2.0's extended capabilities using synthetic scenarios approximating real-world clinical use cases and demonstrate its new capabilities within the context of atrial fibrillation.

Keywords: Clinical practice guidelines · Multi-morbidity · Planning

1 Introduction

Clinical practice guidelines (CPGs) are statements developed systematically from available evidence to assist practitioners in the management of a patient with a specific disease and their application improves quality of care and patient outcomes [1]. Yet their adoption in clinical practice is lacking and one of the major obstacles is the limited support for complex patients suffering from multi-morbidity [2]. Disease management of these

patients is multidimensional by nature, and apart from the clinical dimension needs to consider care settings and patient preferences to maximize patient outcomes.

In this work, we present a significant expansion of our previous planning-based approach to mitigation of adverse interactions between recommendations coming from different CPGs [3] that improves support for multi-morbid patient management by enriching the representation of patient and clinical information. This improvement allows us to provide support for shared decision making by explicitly capturing the multidimensional nature of treatment using a multivariate *objective function* customized to a specific patient encounter. We demonstrate how our improved framework, called MitPlan 2.0, builds upon and extends our previous work using both synthetic and clinical examples. We highlight its use across clinical settings, position it within the context of related work, and describe future work to realize our goal of integrating MitPlan 2.0 into a clinical decision support system used at the point of care.

2 MitPlan 2.0

In planning, a planner is given an initial state of the world, a set of desired goals, and a set of planning actions to find a sequence of planning actions that are guaranteed to generate a new state that satisfies the desired goal(s). Each planning action has a set of parameters, preconditions that must be true for the action to be taken, and effects resulting from its execution. These planning actions are also characterized by a duration, conditional effects, and a cost.

The original version of MitPlan (which we refer to as MitPlan 1.0), described in detail in [3], addressed the problem of mitigation by combining an algorithmic approach with the use of a planner. MitPlan 1.0 accepted as input patient data, patient preferences, the length of a planning horizon, and clinical goals and produced a safe management plan (i.e., where all adverse interactions were addressed), executed within the specified planning horizon. It detected adverse interactions using a combination of revision predicates and revision actions in its domain and mitigated these adverse interactions algorithmically and outside the planner. Specific to the revision operators in the MitPlan 1.0, the planner would terminate having identified what actions needed to be replaced/deleted/added to mitigate adverse interactions between different extended Actionable Graphs (AGs) [3] representing CPGs. The MitPlan 1.0 algorithm would then create a revised problem instance with new actions representing the required revisions and the planner would be applied to this instance. The process would be repeated until the management plan inferred by the planner required no further revisions. In addition to disjointed algorithmic and planning requirements for generating management scenarios, MitPlan 1.0 had very basic support for a single cost associated with the plan.

MitPlan 2.0 addresses these shortcomings by taking a fully planning-based approach, bringing the process of mitigation into the planning space via a new encoding of the planning problem in PDDL. We expand the extended AG to include the nodes derived from the underlying CPG and add to it all nodes introduced by revision operators that are possibly applicable to the AG. The extended AG encapsulates the contingencies introduced by revision operators, making them available to the planner much like patient preferences in MitPlan 1.0 [3].

Revision operators come from knowledge repository (KR). The representation of revision operators is flexible to include clinical actions found in many CPGs, clinical actions for a specific CPGs, and a combination of both. When applied to a specific AG, a revision operator is translated to a binary vector of length equal to the number of nodes in the AG, where each vector element indicates if a node is part of an adverse interaction addressed by the given revision operator. In this way a binary vector can represent multiple revision operators each with different costs. An adverse interaction is present only if a chosen path through the extended AG contains *all* the nodes flagged in a binary vector (if only a subset of nodes are contained, the adverse interaction is not present). If detected, the planner is forced to search for an alternate path in the AG, avoiding the adverse interaction. The binary vectors are automatically created for each planning problem instance, making their implementation scalable.

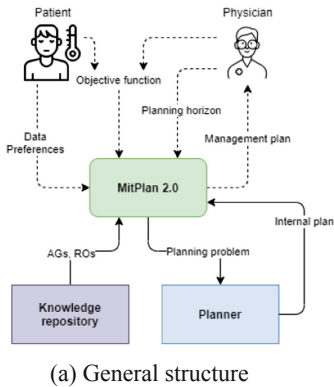
All nodes in the AG are associated with costs, and revision actions are designed to be costlier than the actions they (potentially) revise. The planner prioritizes paths through the AGs with no adverse interactions, finding a clinically feasible plan first and foremost, even if such a plan is more expensive than one that is clinically infeasible. If a path without executing a revision action does not exist, the planner chooses an alternate path containing revision actions that mitigate the adverse interactions and reach the goal nodes with no adverse interactions present. Because revision information is already built into the extended AG, the planner can optimize over various alternative paths, selecting the path associated with the lowest cost. As a simplification, MitPlan 2.0 implements cost minimization however utilities and their maximization are easily supported.

A revision action may itself introduce new adverse interactions, resulting in second-order adverse interactions. MitPlan 1.0 processed revisions sequentially and the plan depended on the order of revisions applied. In contrast, MitPlan 2.0 optimizes over all revision information in a single run and returns the optimal plan, if one exists, selecting the order of revisions to optimize the defined objective function.

MitPlan 1.0 treated patient preferences (e.g., preferred way of drug administration) as alternatives nodes in an AG [3] that are selected based on their dispreference costs (i.e., lower cost indicates a higher preference level). MitPlan 2.0 presents a natural extension of this approach, by unifying patient preferences and revisions, and modelling both as cost-based alternatives. At the same time, it considers more than one metric, each corresponding to a certain dimension, when looking for an optimal path. MitPlan 2.0 employs an objective function given as a weighted sum of selected cost metrics where weights indicate the importance of specific metrics. Possible metrics include various clinical (resources, specialists, capacity, etc.) and patient (financial, burden, preferences, adherence, etc.) indicators, thus it is possible to specify objective functions for a wide range of care settings (e.g., urban or rural), patient populations with unique attributes (e.g., health literacy, income, attitude towards treatment, etc.), and care planning approaches (physician-, nurse-, and patient-centered).

The overview of MitPlan 2.0 is given in Fig. 1(a) and the pseudocode illustrating its operation is given in Fig. 1(b). MitPlan 2.0 is invoked for a specific patient/physician encounter and takes patient data, patient preferences, a planning horizon, and an objective function as input. The objective function can be defined by both the physician and the

patient in a shared decision-making process to combine patient- and the provider-oriented perspectives of management.



input : data: patient data, prefs: patient preferences, horizon: planning horizon, objFunc: objective function for plan optimization
output: managementPlan: management plan

- 1 AGs := select from **knowledge repository** extended AGs representing CPGs used to manage the patient
- 2 ROs := select from **knowledge repository** revision operators possibly applicable to AGs
- 3 problem := create a planning problem using AGs, ROs, data, prefs, horizon and objFunc
- 4 internalPlan := apply **planner** to problem to find an optimal plan
- 5 managementPlan := post-process internalPlan to management plan
- 6 return managementPlan

(b) Operations (pseudo-code)

Fig. 1. Overview of MitPlan 2.0

MitPlan 2.0 starts by retrieving from the KR extended AGs that represent CPGs applied to manage conditions of the patient. It then identifies and retrieves from the KR revision operators (ROs) that are applicable to AGs selected in the previous step (see [3] for a more detailed explanation). Subsequently, it creates a planning problem in PDDL based on extracted AGs and ROs and provided input (patient data and preferences, planning horizon, and objective function). A planner solves the planning problem to obtain an internal plan optimized according to the provided objective function. Some of the actions in the internal plan are then filtered out as they represent implementation details of the planner (e.g., reaching a goal node) and as such are not relevant for patient management (e.g., performing a test, prescribing a drug). The management plan is finally presented to the physician and the patient.

MitPlan 2.0 significantly changed the management of revisions as stated above and the algorithmic part is no longer needed. We note that clinical quality of the final management plan depends on the completeness of the KR and the quality of available revision operators. If the KR is incomplete and does not contain an operator(s) addressing a certain interaction, then the resulting plan may be clinically unsafe. Moreover, if the KR contains a revision operator for a specific interaction that is not based on the most recent evidence, then the obtained plan may be clinically sub-optimal even though it optimized the objective function. The maintenance of the KR in knowledge driven CDSSs is an important part of their life cycle [4]. While it is beyond the scope of this work, we acknowledge it is critical for the clinical validity of MitPlan 2.0.

3 Illustrative Example

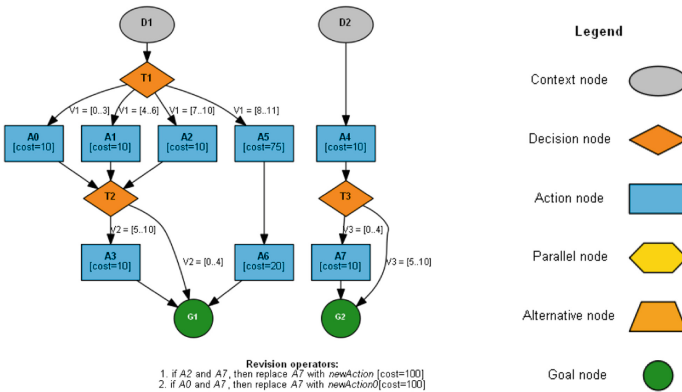
We illustrate the extensions introduced in MitPlan 2.0 through examples. We use synthetically generated examples, grounded in real-world applications and vetted by the

physician on our team, to highlight how the embedding of costs (Sect. 3.1) and the new representation of revisions (Sect. 3.2) in the planning problem expands the capabilities of the planning framework. We use a subset of actions from the Atrial Fibrillation (Afib) CPG (Sect. 3.3) to show the exploration enabled by our new approach.

In each extended AG, Dx represents a context node (i.e., the disease that the AG represents), Ax a clinical action, Tx a clinical test (decision), Vx a patient value, and Gx a goal node representing the successful completion of treatment for the corresponding disease. For the MitPlan 2.0 generated plans, each line lists the planning action taken, the time step it is taken in at the start, and its duration at the end.

3.1 Cost Optimization

MitPlan 2.0 makes use of an objective function when finding a plan to satisfy all goals, where satisfying all goals means reaching all the goal node in all AGs. In Fig. 2(a) we illustrate the use of a simple objective function with a single cost metric (execution cost). Notice in this example that depending on the values of $V1$ the planner may choose $A2$ or $A5$. Suppose $V1 = 9$ and $V3 = 1$. While $A5$ is costlier than $A2$, choosing $A2$ will cause an adverse interaction and necessitate a revision. Because revisions are assigned a higher execution cost than original actions in the AGs (i.e., $cost(A7) < cost(newAction)$), the planner opts to traverse $A5$ and $A6$ to reach goal $G1$ to minimize the overall cost of the plan. Figure 2(b) shows the generated plan.



(a) AGs and available revisions

```


0.000: (makefirstdecision d1 t1 a5) [0.000]
0.001: (takefirstaction d2 a4 t3) [0.000]
0.002: (takeactiontoactionnode d1 a5 a6) [0.000]
0.002: (makedecisiontonode d2 t3 a7) [0.000]
0.004: (takeactiontogoal d2 a7 g2) [0.000]
0.005: (takeactiontogoal d1 a6 g1) [0.000]
0.006: (finalgoalreached d2 g2) [0.001]
0.006: (finalgoalreached d1 g1) [0.001]
    
```

(b) Internal plan generated by MitPlan 2.0

Fig. 2. Selection of minimum cost path

3.2 Revision Application

We demonstrated in Sect. 3.1 how execution costs of an action and optimization come into play when selecting actions. Here we build on that example to show how execution costs are used to select and prioritize revisions when adverse interactions are present. Consider the extended AG in Fig. 3(a) with patient values $V1 = 7$ and $V3 = 3$. In this case, an adverse interaction is unavoidable and one of the two available revisions must be taken. The two choices for replacing $A7$ are *newAction*, which costs 100 units, and *newAction2*, which costs 50 units. Since *newAction2* is less expensive, the planner prioritizes *newAction2*, replacing $A7$ with it. The generated plan is shown in Fig. 3(b). In MitPlan 1.0, the prioritization of revisions would have been done algorithmically and outside the planning process by revising a problem instance and passing it to the planner. Now the selection *and* application of revisions is embedded in the planner’s search for an optimal plan.



- Revision operators:**
 1. if A2 and A7, then replace A7 with *newAction* [cost=100]
 2. if A2 and A7, then replace A7 with *newAction2* [cost=50]

(a) AGs and available revisions

```

0.000: (makefirstdecision d1 t1 a2) [0.000]
0.001: (takefirstaction d2 a4 t3) [0.000]
0.002: (makedecisiontonode d2 t3 newaction2) [0.000]
0.002: (takeactiontodecisionnode d1 a2 t2) [0.000]
0.003: (makedecisiontonode d1 t2 a3) [0.000]
0.004: (takeactiontogoal d2 newaction2 g2) [0.000]
0.005: (takeactiontogoal d1 a3 g1) [0.000]
0.006: (finalgoalreached d2 g2) [0.001]
0.006: (finalgoalreached d1 g1) [0.001]
    
```

(b) Internal plan generated by MitPlan 2.0

Fig. 3. Prioritization of revisions

We use the extended AG in Fig. 3(a) also to demonstrate how MitPlan 2.0 finds an optimal plan when multiple revisions are required (the second revision operator

changed to *if A4 and A3 then replace A3 with newAction1* [$cost = 100$] for the sake of this example). In this problem instance, the patient values for V1 and V2 fall within the range $[5 \dots 10]$ and V3 falls within the range $[0 \dots 4]$, thus necessitating both revisions. Figure 4 shows the resulting plan that does just that. This case demonstrates that MitPlan 2.0 can apply any number of revisions necessary to find an optimal plan while tracking and accounting for any second-order effects from the applied revisions. This is a significant improvement over MitPlan 1.0 where each revision was applied individually by revising the problem instance and rerunning the planner.

```

0.000: (makefirstdecision d1 t1 a2) [0.000]
0.001: (takefirstaction d2 a4 t3) [0.000]
0.002: (makedecisiontonode d2 t3 newaction) [0.000]
0.002: (takeactiontodecisionnode d1 a2 t2) [0.000]
0.003: (makedecisiontonode d1 t2 newaction1) [0.000]
0.004: (takeactiontogoal d2 newaction g2) [0.000]
0.005: (takeactiontogoal d1 newaction1 g1) [0.000]
0.006: (finalgoalreached d2 g2) [0.001]
0.006: (finalgoalreached d1 g1) [0.001]

```

Fig. 4. Internal plan generated by MitPlan 2.0 when multiple revisions are required

3.3 Clinical Illustrative Example

At the core of our mitigation framework is the incorporation of external knowledge for mitigating adverse interactions not found in individual CPGs. By including revisions within the planning space (example in Sect. 3.2), MitPlan 2.0 is able to keep track of the application of these mitigating actions, their sequence of application, and their second-order effects. When clinical mitigation strategies and their effects are embedded into the planning process and considered within an objective function, the resulting management plan is constructed by holistically considering the impact mitigation strategies have on patient management. Iteratively mitigating adverse interactions, as done in the MitPlan 1.0 fails to account for second-order effects.

To demonstrate the power of using an objective function involving multiple metrics, consider the subset of the Afib AG, generated from the CPG, shown in Fig. 5 (full AG is presented in [3] and not reproduced here due to space limitations). In this example, an Afib patient with CHA_2SDS_2 -VASc score greater or equal to 1 can be prescribed an anticoagulant such as Warfarin (*WARF*, dosage 5 mg daily) or one of the direct oral anticoagulants (DOACs) such as Dabigatram (*DABI*, dosage 110 or 150 mg twice daily). For a CHA_2SDS_2 -VASc score less than 1, they are prescribed low dose aspirin (*ASP*).

MitPlan 2.0 can consider financial cost, patient’s burden, or the “cost” of adherence, amongst others, during planning to generate different management scenarios. Let’s first consider each independently. **Financial Cost:** the annual cost of anticoagulation treatment (at a dosage mentioned above) with Warfarin is about \$55USD and with Dabigatram is about \$1200USD. We therefore assign a financial cost of 55 to the WARF node and 1200 to the DABI node in the AG. When financial cost minimization is the optimization goal, treatment with Warfarin is the returned option (see the plan in Fig. 6(a)).




Fig. 5. Subset of the Afib AG from [3]

```
0.000: (makefirstdecision d cha2sds2 altx12) [0.000]
0.001: (makedecisiontonode d altx12 warf) [0.000]
0.002: (takeactiontogoal d warf g) [0.000]
0.003: (finalgoalreached d g) [0.001]
```

(a) Financial cost optimized plan

```
0.000: (makefirstdecision d cha2sds2 altx12) [0.000]
0.001: (makedecisiontonode d altx12 dabi) [0.000]
0.002: (takeactiontogoal d dabi g) [0.000]
0.003: (finalgoalreached d g) [0.001]
```

(b) Weighted cost optimized plan

Fig. 6. Afib generated plans for various cost considerations

Patient’s Burden: A patient on Warfarin needs their international normalized ratio (INR) value checked as a stand-in to standardize the results of the prothrombin time test. Checking the INR value requires a blood test. When a patient is started on Warfarin, this test needs to be done weekly. When their INR value stabilizes, the frequency of the test falls between 2–4 weeks. Performing each blood test requires the patient to visit a laboratory service, imposing an additional treatment burden. On the other hand, a patient on Dabigatram (or any other DOAC) does not need a blood test to measure their INR value and does not incur any additional treatment burden. Using MitPlan 2.0, we assign a burden cost to WARF that is greater than the burden cost of DABI. When minimizing the burden is the optimization goal, Dabigatram is the action returned by MitPlan 2.0.

Adherence Likelihood: It is well documented in literature [5], that Afib patients on anticoagulation treatment poorly adhere to their prescribed medication. Considering that a patient on Dabigatram only takes pills and does not need to get a blood test, their only checkpoint for adherence is during an annual visit with their specialist. On the other hand, a patient on Warfarin needs to have regular blood tests in order to measure that their INR value is in the optimal range and any adjustment of medication dosage requires a visit with a specialist. Consequently, adherence to treatment with Warfarin is typically higher than treatment with Dabigatram due to regular consults with a specialist. As such, MitPlan 2.0 assigns a lower adherence cost to Warfarin than to Dabigatram. When the minimization of non-adherence is the desired goal, regardless of a patient’s burden consideration, treatment with Warfarin is the returned option.

Multiple Weighted Considerations: Typically, there are multiple treatment options for multimorbidity. The selection of a given option is made by a physician in consultation with the patient. MitPlan 2.0 uses a complex objective function to represent multiple dimensions such as a patient’s burden, financial cost or perceived medication adherence and weighs each one differently based on the information gathered during the shared decision-making process. These weights can be revised and adjusted to reflect changing patient’s attitude and evolving clinical context. In Fig. 6(b) we show the plan where minimizing a patient’s treatment burden is considered more important than both the financial cost of a treatment and the perceived patient’s adherence using the multivariate objective function $0.2 * cost + 0.6 * burden + 0.2 * nonadherence$ (all weights are rescaled to the range [0, 1.0]). Note how this plan differs from the one in Fig. 6(a) when only financial cost is considered. The ability to customize a multivariate objective function to a specific patient and encounter represents a powerful tool supporting shared decision-making.

4 Discussion and Future Work

MitPlan 2.0 fully encapsulates the mitigation problem within a planning context with a unified approach to supporting preferences and revisions of CPGs’ recommendations, simultaneously handling multiple revision operations, and optimizing across different metrics. Using various costs associated with nodes in the extended AG, MitPlan 2.0 finds management plans that are optimized according to a weighted multivariate objective function. This approach supports the combination of clinical dimensions with different treatment aspects such as financial cost, patient’s burden, patient’s perceived adherence to treatment, or cost of clinical resources required for treating a patient.

There are other approaches to mitigation of adverse interactions among multiple CPGs reported in the literature and below we briefly summarize the ones that are most closely related to MitPlan 2.0. META-GLARE [6] considers temporal characteristics of CPG actions during mitigation, employs goals to control the planning process, and it has been extended to model physician preferences [7]. However, no different types of preferences or costs are considered and there is no optimization over possible plans. Jafarpour et al. [8] propose an ontology-based framework for integrating multiple CPGs during execution time using policies. Some of these policies can optimize the use of clinical resources but other types of costs or preferences are not considered and there is no global optimization of generated plans.

A multi-agent planning (MAP) framework [9] is used to automatically generate several candidate management plans, evaluate them according to predefined patient- and institution-related metrics, and select the optimal one that minimizes the overall cost. MAP requires a more complex computational framework with multiple agents, and it assumes that secondary knowledge representing revision operators is embedded in CPGs. This assumption makes it more difficult to maintain than MitPlan 2.0, which uses an external knowledge repository for the revision operators. Finally, Kogan et al. [10] propose a goal-driven mitigation framework that uses standard representations (PROforma and HL7 FHIR), relies on existing knowledge sources, operates on different levels of

abstraction, and generates explanations for proposed mitigations. Yet, it does not consider patient preferences and other costs to prioritize candidate plans as the final selection of a management plan is made by the clinician.

We are pursuing several directions for future work. The use of hard and soft constraints, that is, constraints that must be satisfied in any solution (e.g., a hypertension drug must be administered) and those that are optional (e.g., a blood test applied only if lab resources are available), respectfully, will add an additional level of personalization. We are also exploring how to generate management plans from a partial satisfaction of a subset of defined goals. Finally, we are studying the use of a stochastic model of mitigation. As the representational and functional complexity of the CPGs and in turn the planning problem increases, it will be necessary to shift from a deterministic representation towards one that supports probabilities tied to the execution of actions. In the real-world, medication may not be taken by a patient or may not have the intended effects, test results may be inconclusive or inaccurate, and future test results assumed for the sake of planning could return unanticipated values.

Acknowledgements. We thank Andrew Coles and Amanda Coles for their clarifications regarding PDDL and OPTIC and the reviewers for their helpful comments.

References

1. Goud, R., et al.: The effect of computerized decision support on barriers to guideline implementation: a qualitative study in outpatient cardiac rehabilitation. *Int. J. Med. Inform.* **79**(6), 430–437 (2010)
2. Peleg, M.: Computer-interpretable clinical guidelines: a methodological review. *J. Biomed. Inform.* **46**(4), 744–763 (2013)
3. Michalowski, M., Wilk, S., Michalowski, W., Carrier, M.: MitPlan: a planning approach to mitigating concurrently applied clinical practice guidelines. *Artif. Intell. Med.* **112**, 102002 (2021)
4. Sutton, R.T., Pincock, D., Baumgart, D.C., et al.: An overview of clinical decision support systems: benefits, risks, and strategies for success. *npj Digit. Med.* **3**, 17 (2020)
5. Raparelli, V., Proietti, M., Cangemi, R., Lip, G.Y., Lane, D.A., Basili, S.: Adherence to oral anticoagulant therapy in patients with atrial fibrillation. Focus on non-vitamin K antagonist oral anticoagulants. *Thromb. Haemost.* **117**(2), 209–218 (2017)
6. Bottrighi, A., Piovesan, L., Terenziani, P.: Coping with “exceptional” patients in META-GLARE. In: Cliquet Jr., A., et al. (eds.) *BIOSTEC 2018*. CCIS, vol. 1024, pp. 298–325. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29196-9_16
7. Terenziani, P., Andolina, A.: Considering temporal preferences and probabilities in guideline interaction analysis. In: Riaño, D., Wilk, S., ten Teije, A. (eds.) *AIME 2019*. LNCS (LNAI), vol. 11526, pp. 120–124. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-21642-9_16
8. Jafarpour, B., Abidi, S.R., Woensel, W.V., Abidi, S.S.R.: Execution-time integration of clinical practice guidelines to provide decision support for comorbid conditions. *Artif. Intell. Med.* **94**, 117–137 (2019)

9. Fdez-Olivares, J., Onaindia, E., Castillo, L., Jordan, J., Cozar, J.: Personalized conciliation of clinical guidelines for comorbid patients through multi-agent planning. In: Proceedings of Artificial Intelligence in Medicine (AIME 2019), pp. 167–186 (2019)
10. Kogan, A., Peleg, M., Tu, S.W., Allon, R., Khaitov, N., Hochberg, I.: Towards a goal-oriented methodology for clinical-guideline-based management recommendations for patients with multimorbidity: GoCom and its preliminary evaluation. *J. Biomed. Inform.*, **112**, 103587 (2020)