

Generalizing the Detection of Internal and External Interactions in Clinical Guidelines

Veruska Zamborlini^{1,2}, Rinke Hoekstra^{1,3}, Marcos da Silveira², Cedric Pruski², Annette ten Teije¹ and Frank van Harmelen¹

¹Dept. of Computer Science, VU University Amsterdam, Amsterdam, The Netherlands

²LIST Luxembourg Institute of Science and Technology, Esch-sur-Alzette, Luxembourg

³Faculty of Law, University of Amsterdam, Amsterdam, The Netherlands

Keywords: Clinical Guidelines, Semantic Web, Knowledge Representation, Ontologies.

Abstract: This paper presents a method for formally representing Computer-Interpretable Guidelines to deal with multimorbidity. Although some approaches for merging guidelines exist, improvements are still required for combining several sources of information and coping with possibly conflicting pieces of evidence coming from clinical studies. Our main contribution is twofold: (i) we provide general models and rules for representing guidelines that expresses evidence as causation beliefs; (ii) we introduce a mechanism to exploit external medical knowledge acquired from Linked Open Data (Drugbank, Sider, DIKB) to detect potential interactions between recommendations. We apply this framework to merge three guidelines (Osteoarthritis, Diabetes, and Hypertension) in order to illustrate the capability of this approach for detecting potential conflicts between guidelines and eventually propose alternatives.

1 INTRODUCTION

Clinical Guidelines (CG) are developed for supporting physicians decision, e.g. specifying what treatment work best in what **situation** (Peleg, 2013). When possible, the **recommendations** provided by CGs are based on **evidence** from clinical researches. In this case, there is a direct mapping to the clinical evidence that describes the effects (**transitions**) of certain **care action** (e.g. *do not administer aspirin because of an increased risk of gastrointestinal bleeding*). Since an evidence is not a fact, a multitude of evidence rating systems (Lohr, 2003) are adopted by CGs authors. Epistemologically, an evidence reflects a **belief** in the existence of a **causal** relation between e.g. *administering aspirin and gastrointestinal bleeding*. Furthermore, CGs are targeted to the treatment of a specific illness. However, it is quite common to have patients with *multiple* diseases (multi-morbidity) that need to be addressed according to different CGs. For example, according to (Barnett et al., 2012), around 40% of 55 years old patients suffer from at least 2 diseases, and 20% of 70 years old patients suffer from at least 4 diseases in Scotland. As with any large volume of regulations, combined guidelines almost

inevitably involve intricate **interactions** between the recommendations they describe. Finding interactions (like potential conflicts) requires intensive collaboration in multidisciplinary teams.

Computational support can be of great value for supporting physicians to handle all this complexity. Many languages have been proposed for representing “computer interpretable” guidelines (CIG) and reasoning about it (Peleg, 2013). However, the concepts here discussed are poorly or not addressed by those approaches. The main reason is because much has been devoted to executing guidelines within treatments rather than other purposes such as combining and updating CGs. In particular, regarding the issue of multimorbidity, existing approaches for combining CGs are limited in their ability to automatically detect the interactions, propose alternatives or combining more than two guidelines (Zamborlini et al., 2015b).

This work follows an incremental methodology. We start by addressing realistic but simplified case studies, and add more complexity according to the lessons learned in each iteration. Therefore, this paper is the continuation of earlier work reported in (Zam-

borlini et al., 2014a; Zamborlini et al., 2014b; Zamborlini et al., 2015b; Zamborlini et al., 2015a). In this series of work, we investigated (i) what knowledge is required to represent and reason about CGs (rather than how to acquire such knowledge), particularly for supporting the multimorbidity issue; (ii) how it can be formalized; (iii) how it can be implemented using Semantic Web technologies, so that (iv) we can exploit the medical knowledge available as Linked Open Data (LOD). As a consequence, the results we obtain are limited to the current expressiveness of the model, e.g. temporal aspects and related interactions will be addressed in future iterations.

This paper reports on improvements to both the models and the implementation to better address the issue of multimorbidity. The contributions are (C1) a more generic version of the models with respect to recommendations, beliefs and event types. This includes (C2) a formalization of the improved models and rules in FOL; and (C3) a Semantic Web framework for representing and reasoning about recommendations and beliefs using standard vocabularies. This provides (C4) a flexible mechanism for reusing external knowledge bases to extend our ability to detect interactions (showcased using DrugBank and Sider).

The remainder of this paper is as follows: Sect. 2 presents a case study to illustrate the main concepts, which are further defined in the models and rules, followed by their implementation. An experimental assessment shows the results obtained for the referred case study in Sect. 3. The related work is discussed in Sect. 4 and the main contributions and future work are discussed in Sect. 5.

2 THE MODELS & FRAMEWORK

This section describes our case study, as well as the adapted version of the TMR (Transition-based Medical Recommendation) models and their semantic web-based implementation as a framework for representing norms (recommendations) in the clinical domain and reasoning about interactions among them.

2.1 Case Study

This case study is meant for illustrating the concepts previously mentioned and further defined in the next section. It concerns the combination of three guidelines, namely Osteoarthritis (OA), Diabetes (DB) and Hypertension (HT), and the detection of interactions among them (adapted from (Jafarpour, 2013)). The recommendations, depicted in Fig. 1, are:

Diabetes (DB)

1. *Should adm. NSAID to reduce blood coagulation*
2. *Should adm. Tramadol to reduce blood coagulation*
3. *Should adm. Insulin to reduce blood sugar level*

Osteoarthritis (OA)

1. *Should NOT administer Aspirin to avoid increasing the risk of gastrointestinal bleeding*
2. *Should administer Ibuprofen to reduce pain*

Hypertension (HT)

1. *Should adm. Thiazide to reduce the blood pressure*

Among them some interactions can be identified:

Internal Interactions:

1. DB.1 and DB.2 are alternative recommendations meant for promoting the same effect.
2. DB.1 and OA.1 are contradictory recommendations since the first might lead to the prescription of *Aspirin* which is non-recommended by the later.

External Interactions (from external knowledge sources):

1. DB.1 and DB.2 both have as external alternative *Administer Epoprostenol* to achieve the desired effect according to *DrugBank*.
2. DB.2 and OA.2 recommend incompatible actions according to *Drugbank*
3. DB.3 and HT.1 interact since the latter might lead to prescription of *bendroflumethiazide*, which has *high blood sugar level* as side effect according to *Sider*, as opposed to goal in the former.
4. HT.1 and OA.2 interact since the latter has *high blood pressure* as side effect according to *Sider*, as opposed to goal in the former.
5. (others)

In Fig. 1 the big rectangles in both left and right sides represent **beliefs** regarding the **care actions** (*administering tramadol*). The latter is represented as dotted ellipses inside the beliefs. The **causation beliefs** are about a **transition** between **situations** (*blood coagulation goes from normal to low*) that are believed to be promoted by executing a care action type). The causation belief has a **frequency** e.g. *administer tramadol always reduce the blood coagulation*. For sake of simplicity, we consider in this work only *always* as frequency for all causation beliefs. They also have a **strength** associated, which corresponds to the evidence level (e.g. *high level*), according to the quality attributed to the **sources** (or studies) that provide such knowledge. The beliefs in gray shade represent the knowledge imported from an external source described in the top left (e.g. *DrugBank*). The strength in this case will depend on the reliability of each data source. The external sources here considered describe two types of beliefs: causation belief or **incompatibility belief**. The latter represents action types that should not be recommended together, e.g. *Administer Aspirin* is incompatible with

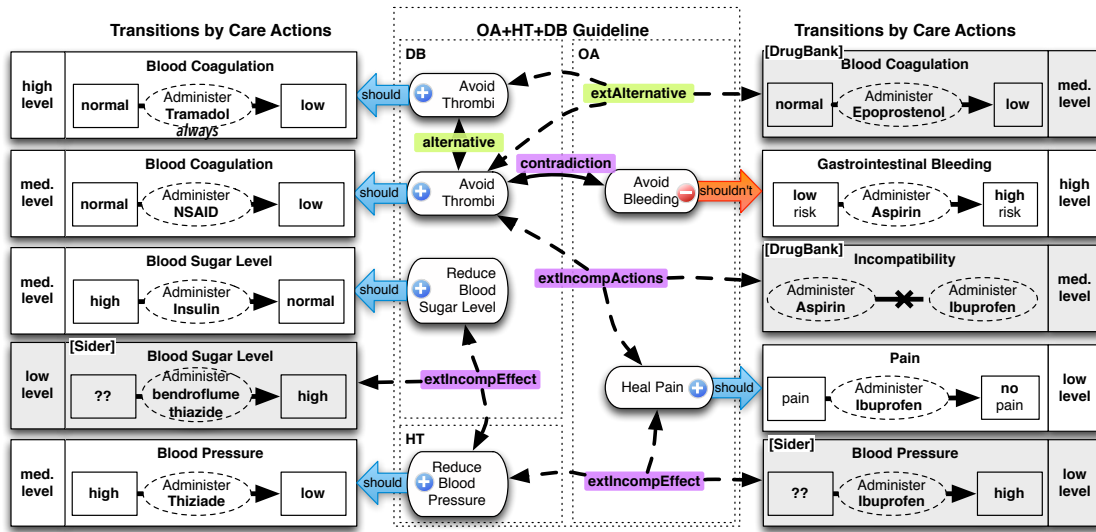


Figure 1: Case study on combining guidelines for OA+HT+DB.

Administer Ibuprofen (the reason is not provided in structured way from the sources).

The dotted rectangles in the middle represent the **guidelines**. The more external one is the merge of the three guidelines for OA+HT+DB. They comprise both the **recommendations** (e.g. *avoid thrombi*) and the **interactions** *alternative* among them. The former is represented as rounded rectangles, and the latter is depicted by labelled thin arrows connecting the interacting recommendations and beliefs. Solid arrows are for internal interactions and dotted arrows for external ones. A positive (or negative) recommendation is indicated by a thick arrow labeled with “should” (or “should not”).

2.2 Conceptual Model & Rules

Figure 2 presents a UML class diagram for the TMR_{Event} model describing some relevant concepts and relations regarding **event types** in the scope of this work. The concepts introduced in previous versions of the models are depicted in gray-shade (same for the next diagrams). This model is inspired in UFO (Unified Foundational Ontology) (Guizzardi et al., 2013) that is a formal theory describing some of the general concepts used here, namely Type (Universal) and Category, as well as Object, Event, Action and Situation Types. The model regards mostly types of things¹ since it is meant for modeling, for example, the type of event that is expected as consequence of another one, rather than the particular event that was the consequence of another particular one. In other words, we do not want to say that John’s pain was re-

¹For sake of simplicity we can omit the word ‘type’.

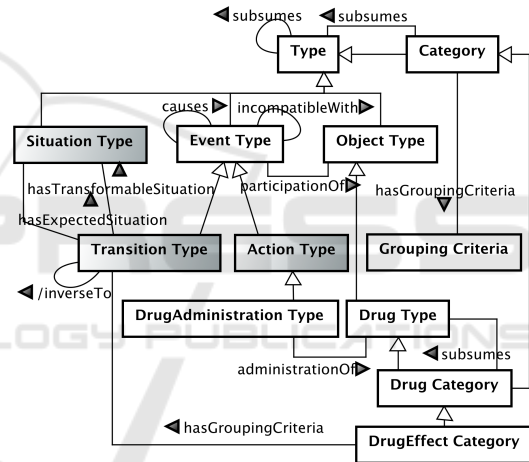


Figure 2: UML class diagram for the TMR_{Event} Model.

lieved due to the administration of aspirin, but that administering aspirin often relieves the pain of patients².

While **action types** concern event types to be performed by an intentional agent (omitted in the model), **transition types** concern (deterministic) event types for which pre and post **situation types** can be defined³. In other words, it represents the transformation of a **situation type** into another (transformable & expected situations). An event type can be defined as the **participation of** a certain **object type**, e.g. a *DrugAdministration type* is the administration of (participation of) a *Drug type*. An event type can **cause** another event type to happen (occurrence of one causes the occurrence of the other). Moreover,

²For a deeper explanation see (Zamborlini et al., 2014a).

³We do consider some event types are non-deterministic or non-intentional, but this is out of scope of this work.

an event type can also be **incompatible with** another one when they can not or should not occur together. In other words, either happening together is not possible or would bring about results/transitions that are not the expected ones.

A **category** is a type that (transitively) subsumes (or regroups) other types according to a **grouping criteria**, e.g. *ThiazideDrug* is a category of drugs that contains the molecule *thiazide*, e.g. *bendroflumethiazide*. In this case the grouping criteria regards a structural property. However, it can also concern the effect expected to be promoted, e.g. *NSAID* is the category for the (non-steroidal) drugs expected to promote the transition *reduce inflammation*, e.g. *Aspirin*.

FOL rules are provided for deriving relations relevant in the context of this work. Some relations are defined in terms of other relations, for example, *inverseTo* between transition types is one transition that ‘undo’ the effect of the other. These relations are preceded by a slash in the models (previous and forthcoming). Other relations can be propagated, e.g. the grouping criteria of an object category (Analgesic Drug) is the grouping criteria of the event category of administering such object (Adm. Analgesic) and *vice-versa*.

Inverse Transitions: one transition type $t1$ transforms situation $s1$ into $s2$ while another transition $t2$ transforms $s2$ into $s1$.

R.1 $\forall t1, t2, s1, s2$ TransitionType($t1$) \wedge TransitionType($t2$)
 \wedge SituationType($s1$) \wedge SituationType($s2$) \wedge $s1 \neq s2$
 \wedge hasTransformableSituation($t1, s1$)
 \wedge hasExpectedSituation($t1, s2$)
 \wedge hasTransformableSituation($t2, s2$)
 \wedge hasExpectedSituation($t2, s1$)
 \rightarrow inverseTo($t1, t2$)

Propagating grouping criteria via administrationOf: if an action type $a1$ is the administration of a drug type $d1$, a transition type $t1$ is the grouping criteria for $d1$ iff it is also for $a1$.

R.2 $\forall a1, d1, t1$ (ActionType($a1$) \wedge DrugType($d1$)
 \wedge TransitionType($t1$) \wedge administrationOf($a1, d1$))
 \rightarrow (hasGroupingCriteria($d1, t1$)
 \leftrightarrow hasGroupingCriteria($a1, t1$))

Other relations can be difficult to be precisely defined either for epistemic or ontological issues. For instance, some event types might not have a precise definition of their consequences, e.g. administering aspirin sometimes relieves the pain, sometimes it does not. We address this issue in the TMR_{Belief} model by representing those relations through beliefs, presented in a UML class diagram in Fig. 3. In this work, **beliefs** allow to represent a ‘degree of truth’ for assertions about **things/entities** according to a source.

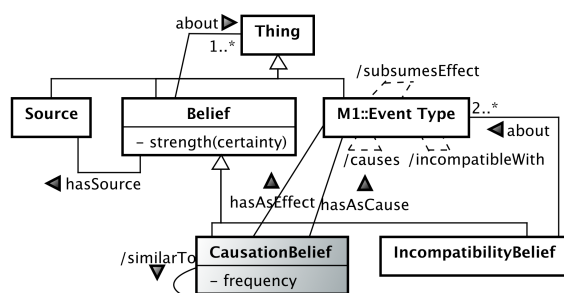


Figure 3: UML class diagram for the TMR_{Belief} Model. ‘M1’ is used as short reference for TMR_{Event} .

It allows for complementary beliefs, e.g. *aspirin relieves the pain in 80% of the cases (therefore it does not in 20%)*, but also for inconsistent ones, e.g. *administering aspirin always relieves the pain and it never relieves the pain*. This is a desired feature since for some assertions there can be no common agreement from different sources (see (Zamborlini et al., 2015a)). However, we consider that one (merged) guideline that rely on incompatible beliefs is then inconsistent. We also account for the certainty/quality of the belief as its **strength**, such as the evidence level classification in clinical guidelines.

We are particularly interested in beliefs about the relations *causes*, *subsumes* and *incompatibleWith* between event types, for which we provide ‘epistemic/doxastic’ versions (represented as dotted lines in the model). In other words, they are relations dependent on the existence of a belief to ground their truthfulness (in practice they have a belief as a third argument). Therefore, they are not the same as the ones in Fig. 2.

The **causation belief** between event types reflects the likelihood/frequency (probability) of one causing the another according to a source. In this work we focus on beliefs about action (*hasAsCause Action Type*) causing transitions (*hasAsEffect Transition Type*) as justification for the clinical recommendations. Moreover, for the interaction rules we consider only the positive causation beliefs, i.e., an action type always cause a transition, since it is not on the scope of this work both (i) the negative causation, which only appears as sub-justifications of recommendation (discussed in (Zamborlini et al., 2015a)) and (ii) the intermediate frequency values (often, rarely, etc)⁴. In its turn, the **incompatibility** between event types is considered in this work to be given as an assertion. Therefore it is represented as a belief, although it could be explained/derived at a certain level of granularity. Fi-

⁴This approach exclude endless assertions about all the effects an event is not expected to produce since the beliefs are defined in CGs or scientific papers by a community of experts, e.g. cancer is not an effect of a certain drug.

the TMR_{Norm} model.

2.2.1 Internal Interaction Rules

Considering the modifications in the model, we propose in this section the corresponding adaptation of the **internal interaction** rules presented in (Zamborlini et al., 2015a) (contribution C2). The following types of interactions are defined:

Repeated Action: two positive norms about the same action or about actions in a subsuming relation. The second rule is for the cumulative behavior, i.e. when two different interactions of this type relate the same norm ($n2$) then those interactions are the same.

(I.1.1) $\forall r, n1, n2, a1, a2, cb1, cb2$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{regulates}(r, n2, a2, \text{'should'}, cb2)$
 $\wedge (a1 = a2 \vee \text{subsumes}(a1, a2) \vee \text{subsumes}(a2, a1))$
 $\rightarrow \exists i(\text{RepeatedAction}(i) \wedge \text{relates}(i, n1)$
 $\wedge \text{relates}(i, n2))$

(I.1.2) $\forall i1, i2, n1, n2, n3$ (RepeatedAction($i1$)
 $\wedge \text{RepeatedAction}(i2) \wedge \text{relates}(i1, n1)$
 $\wedge \text{relates}(i1, n2) \wedge \text{relates}(i2, n2) \wedge \text{relates}(i2, n3)$
 $\wedge r1 \neq n3 \wedge n1 \neq n2 \wedge n2 \neq n3$)
 $\rightarrow i1 = i2$

Alternative Actions: two positive norms about different actions for achieving the same transition, i.e. they are based on similar causation beliefs for different actions. The second rule is for the cumulative behavior

(I.2.1) $\forall r, n1, n2, a1, a2, cb1, cb2$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{regulates}(r, n2, a2, \text{'should'}, cb2)$
 $\wedge \text{similarTo}(cb1, cb2) \wedge a1 \neq a2$
 $\rightarrow \exists i(\text{AlternativeActions}(i) \wedge \text{relates}(i, n1)$
 $\wedge \text{relates}(i, n2))$

(I.2.2) $\forall i1, i2, n1, n2, n3$ (AlternativeActions($i1$)
 $\wedge \text{AlternativeActions}(i2) \wedge \text{relates}(i1, n1)$
 $\wedge \text{relates}(i1, n2) \wedge \text{relates}(i2, n2) \wedge \text{relates}(i2, n3)$
 $\wedge n1 \neq n3 \wedge n1 \neq n2 \wedge n2 \neq n3$)
 $\rightarrow i1 = i2$

Contradictory Norms: (i) two norms, positive and negative, about the execution of same action (or actions in a subsuming relationship) or (ii) two norms, positive and negative, about different actions promoting the same transition or (iii) two positive regulations about different actions for achieving inverse transitions.

(I.3) $\forall r, n1, n2, a1, a2, cb1, cb2, t1, t2$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{regulates}(r, n2, a2, \text{'should'}, cb2)$
 $\wedge \text{causes}(a1, t1, \text{'always'}, cb1)$
 $\wedge \text{causes}(a2, t2, \text{'always'}, cb2)$

$\wedge ((\text{str} = \text{'should-not'}$
 $\wedge (a1 = a2 \vee \text{subsumes}(a1, a2) \vee \text{subsumes}(a2, a1))$
 $\vee (\text{str} = \text{'should-not'} \wedge a1 \neq a2 \wedge t1 = t2)$
 $\vee (\text{str} = \text{'should'} \wedge a1 \neq a2 \wedge \text{inverseTo}(t1, t2)))$
 $\rightarrow \exists i(\text{Contradiction}(i) \wedge \text{relates}(i, n1) \wedge \text{relates}(i, n2))$

Repairable Transition: two norms, positive and negative, about different actions that are believed to cause inverse transitions, i.e. if the undesired effect cannot be avoided, it can be repaired by another action.

(I.4) $\forall r, n1, n2, a1, a2, cb1, cb2, t1, t2$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{regulates}(r, n2, a2, \text{'should-not'}, cb2)$
 $\wedge \text{causes}(a1, t1, \text{'always'}, cb1)$
 $\wedge \text{causes}(a2, t2, \text{'always'}, cb2)$
 $\wedge a1 \neq a2 \wedge \text{inverseTo}(t1, t2)$
 $\rightarrow \exists i(\text{RepairableAction}(i) \wedge \text{relates}(i, n1)$
 $\wedge \text{relates}(i, n2))$

2.2.2 External Interaction Rules

Beliefs from other sources provide interesting information to enrich the system, allowing for detection of **external interactions** defined as:

External-Alternative Actions: actions of which the regulation system is not aware that might promote a desired effect according to external sources. The second rule is for the cumulative behavior, i.e. two different interactions of this type relating the same external causation belief are the same.

(E.1.1) $\forall r, n1, a1, a2, cb1, cb2$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{similarTo}(cb1, cb2) \wedge \text{hasAsCause}(cb2, a2)$
 $\wedge \neg(\exists n2 \text{regulates}(r, n2, a2, \text{'should'}, cb2))$
 $\wedge \neg \text{subsumes}(a1, a2)$
 $\rightarrow \exists i(\text{ExternalAlternativeAction}(i)$
 $\wedge \text{relates}(i, n1) \wedge \text{relates}(i, cb2) \wedge \text{relates}(i, a2))$

(E.1.2) $\forall i1, i2, r, n1, n2, cb$ (ExternalAlternativeAction($i1$)
 $\wedge \text{ExternalAlternativeAction}(i2) \wedge \text{CausationBelief}(cb)$
 $\wedge \text{relates}(i1, cb) \wedge \text{relates}(i1, n1)$
 $\wedge \text{relates}(i2, cb) \wedge \text{relates}(i2, n2)$
 $\wedge \text{Regulation}(r) \wedge \text{partOf}(n1, r) \wedge \text{partOf}(n2, r)$
 $\rightarrow i1 = i2$

External-Incompatible Actions: two positive norms about actions that the regulation system is not aware of them being incompatible to each other (or to a subsumed action), according to external sources.

(E.2) $\forall r, n1, n2, a, a1, a2, cb1, cb2, ib1$
 $\text{regulates}(r, n1, a1, \text{'should'}, cb1)$
 $\wedge \text{regulates}(r, n2, a2, \text{'should'}, cb2)$
 $\wedge \text{incompatibleWith}(a, a2, ib1)$
 $\wedge (a1 = a \vee \text{subsumes}(a1, a) \vee \text{subsumes}(a, a1))$

$\rightarrow \exists i(\mathbf{ExternalIncompatibleActions}(i)$
 $\wedge \text{relates}(i,n1) \wedge \text{relates}(i,n2) \wedge \text{relates}(i,ib11))$

External-Incompatible Effects: effects (situations) recommended to be either avoided or changed, but which are believed to be promoted by other recommended actions, according to external sources, while the regulation system is not aware of it.

(E.3) $\forall r,n1,n2,a,a1,a2,cb,cb1,cb2,s1($
 $\text{regulates}(r,n1,a1,st,cb1)$
 $\wedge \text{causes}(a1,t1,'always',cb1)$
 $\wedge ((st= 'should') \wedge \text{hasTransformableSituation}(t1,s1))$
 $\vee (st= 'should-not') \wedge \text{hasExpectedSituation}(t1,s1))$
 $\wedge \text{causes}(a,t,'always',cb) \wedge a \neq a1$
 $\wedge \text{hasExpectedSituation}(t,s1)$
 $\wedge \text{regulates}(r,n2,a2,'should',cb2) \wedge cb \neq cb2$
 $\wedge (a2 = a \vee \text{subsumes}(a,a2) \vee \text{subsumes}(a2,a))$
 $\rightarrow \exists i(\mathbf{ExternalIncompatibleEffects}(i) \wedge \text{relates}(i,n1)$
 $\wedge \text{relates}(i,n2) \wedge \text{relates}(i,s1) \wedge \text{relates}(i,cb2))$

This section concludes the contribution C2 by providing *generic* rules for detecting external interactions. In (Zamborlini et al., 2015b) the rules were specific for a dataset (namely DrugBank). Now the generic rules apply to beliefs imported from any dataset. At this point its also important to observe that both models and rules are defined in a domain-independent way. A SemWeb-based implementation is provided in the next section.

2.3 SemWeb Implementation

This section presents a SemWeb implementation⁵ for the proposed approach. The proposed models have a straightforward mapping to OWL2 (omitted in the paper). However, for instantiating the models with the clinical knowledge, we propose the use of a framework as RDF graph structure based on the open formats Nanopublication, Provenance and Open Annotation (see section 4). It is applied for representing the recommendations (norms) and beliefs as assertions connected to their sources, besides other meta-information. In the sequence, the implementation of the FOL rules using SWI-Prolog is exemplified, together with a procedure adopted to import clinical knowledge from LOD (e.g. DrugBank and Sider) as assertions via generic predicates (beliefs).

2.3.1 Framework

The framework, illustrated in Fig. 5, follows the *Nanopublication* structure, which presupposes the use of *Prov* vocabulary, and is enriched with (optional) *Open Annotation* vocabulary. The latter is meant for

⁵ Accessible at <http://rapgmsbgym.github.io>.

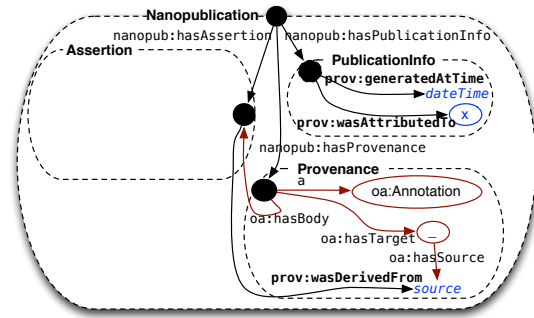


Figure 5: Nanopublication Schema proposed for representing Beliefs & Recommendations.

representing assertions that are (somehow) extracted from textual documents.

Rounded-dotted boxes represent named graphs containing triples, which in turn are represented as directed-named arrows among resources. The black circles represent the named graphs themselves as subject/object, while the other resources are represented as ellipses with a description inside (where underline stands for blank nodes) or as an expected data-value (e.g. source or date).

The more external named graph, called *Nanopublication*, connects the following three named graphs: The *assertion* is a named graph where some knowledge is described using suitable vocabularies, in our case the TMR models. The other named graphs are meant to provide the meta-information about both the assertion and its publication as rdf-data: (i) the *provenance* graph can contain information such as the source (*prov:wasDerivedFrom*) of the assertion (e.g. clinical guideline, study or dataset), and text-annotations (*oa:Annotation*) when the assertion is extracted from a piece of text; and (ii) the *publication-Info* graph provides meta-information such as when the publication was created and by whom. For sake of readability, henceforth we omit from the figures part of the framework that is not relevant for the discussion here conducted.

Figure 6 illustrates the representation of both a causation belief (at the top) and a recommendation (at the bottom). The *Assertion1* is a *CausationBelief* with *high* strength level, about the action type that *Adm. Ibuprofen* always causes the transition type *Pain relief* according to *OA-CIG-Description*. The *Assertion2* is a *Norm*, part of *OA-GIC* that states *Adm. Ibuprofen should* be executed based on the *evidence* stated in *Assertion1* according to *OA-CIG-Description*. Beliefs taken from external sources are similarly represented (see Fig. 7).

The framework favor data reusability as LOD,

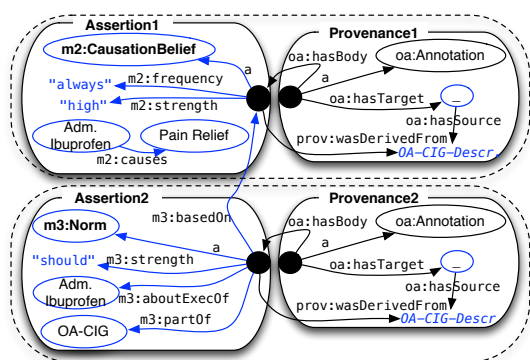


Figure 6: Nanopublication Schema for Causation Beliefs & Norms extracted from Guidelines. ‘M2’ and ‘M3’ are used as prefixes for TMR_{Belief} and TMR_{Norm} respectively.

since it is compatible with SemWeb standards proposed for expressing and annotating knowledge extracted from (scientific) publications. It comprises part of contribution C3 (SemWeb implementation of model and rules).

2.3.2 Rules

The proposed FOL rules have the typical format of Prolog rules, what makes its implementation very straightforward. The implemented rules are here illustrated as: function F.1 implements R.8 while functions F.2.1, F.2.2 implement the rules I.1.1 and I.1.2 for interaction *RepeatedAction*. For the purpose of this application, the existential quantifier in the consequent of interaction rules is implemented as a Prolog function called *existsInteraction*. This function uses the *rdf_assertion* built-in-function to insert the respective interaction in the dataset in case it does not exist.

(F.1) regulates(Reg, Norm, ActT, Str, CBelief) :-
instanceOf(Norm, m3:'Norm'),
rdf(Norm, m3:'partOf', Reg),
rdf(Norm, m3:'aboutExecutionOf', ActT),
rdf(Norm, m3:'strength', literal(type(xsd:string,Str))),
rdf(Norm, m3:'basedOn', Belief, Norm).

(F.2.1) forall((regulation(Reg),
regulates(Reg, N1, ActionT1, 'should', _),
regulates(Reg, N2, ActionT2, 'should', _),
different(N1,N2),
(same(ActionT1, ActionT2)
; subsumes(ActionT1, ActionT2)
; subsumes(ActionT2, ActionT1))),
existsInteraction('RepeatedAction', N1, N2)).

(F.2.2) forall((interacts('RepeatedAction', N1, N2, I1),
interacts('RepeatedAction', N2, N3, I2),
different(N1,N3), different(I1, I2)),
rdf_assert(I1, owl:sameAs, I2)).

This section comprises part of contribution C4 (flexible mechanism for reusing LOD to detect in-

teractions) and together with the framework, it concludes contribution C3.

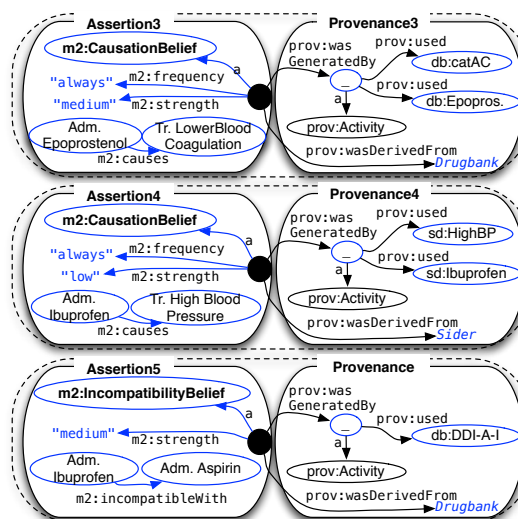


Figure 7: Nanopublication schema for representing beliefs extracted from external knowledge sources. ‘M2’ is used as prefix for TMR_{Belief} .

2.3.3 Using External Knowledge - LOD

For some external knowledge-bases available as LOD, a procedure is adopted for importing the knowledge as new beliefs. For example, knowledge about the (side) effects of administering a certain drug can be imported from *Sider*. Each LOD requires a procedure of reinterpreting the original knowledge into the format here proposed. Once it is done, the rules can derive external interactions regardless to the source⁶. The following SWI-Prolog rule is used for importing causation beliefs from *drugbank* (other rules can be found online, also comprising *sider* and *dikb*).

Drug Alternatives for all drugs belonging to a *drugCategory* regarding an effect, the causation beliefs are asserted about the actions of administering those drugs promoting the referred effect/transition.

```
forall( rdf(DrugCat, model:'hasGroupingCriteria',
Trans1),
same(DrugCat, DrugCatDB),
rdf(DrugDB, drugcategory:'category', DrugCatDB),
rdf(Act1, model:'administrationOf', DrugType),
same(DrugType, DrugDB)),
(assertCausation(Act1, Trans1, 'always', 'drugbank',
```

⁶The Drug and Situation Types are mirrored and mapped to the to the external knowledge sources via owl:sameAs.


```
NanopubURI),
assertProvResourceUsed(NanopubURI, DrugDB),
assertProvResourceUsed(NanopubURI, DrugCatDB))).
```

Figure 7 depicts the representation of beliefs extracted from both Drugbank and Sider. The *Assertion3* at the top is a *CausationBelief* with *medium* strength level, stating that the action type *Adm. Epoprostenol* always causes the transition type *Lower Blood Coagulation* according to *Drugbank*. The *Assertion4* in the middle is another *CausationBelief* with *low* strength level, stating that *Adm. Ibuprofen* always causes to *Higher Blood Pressure* according to *Sider*. Finally, *Assertion5* is a *IncompatibilityBelief* with *medium* strength level, stating that *Adm. Ibuprofen* and *Adm. Aspirin* are incompatible according to *Drugbank*. For all of them, the provenance graphs contain, besides the source dataset, the external resources based on which the assertions were generated.

This section concludes contribution C4 together with the formalisation and implementation of rules for external interactions. Medical guidelines as well as external clinical knowledge can be expressed by means of the conceptual model and can be implemented in a SemWeb-based Framework for automatically detecting interactions. In the next section we provide an experimental assessment by discussing the implementation a case study on detecting recommendations interactions enriched by external knowledge sources.

3 EXPERIMENTAL ASSESSMENT

This section presents the results obtained by implementing the case study on combining OA+HT+DB guidelines (Sect. 2.1). The following activities were performed in the experiment: (i) the guideline knowledge was (manually) introduced in a RDF dataset according to the implementation here proposed; (ii) the rules for importing LOD were fired; (iii) the rules for inferring relations and interactions were fired. An ‘interactive’ documentation describing the experiment and the prolog code is available online⁷. Figure 8 summarizes the obtained results. It describes the type of interaction, the interaction and its source (derived from internal or external knowledge). The first six lines are the interactions described in Sect. 2.1. The last two lines illustrate that more external interactions can be detected, actually much more given the large volume of clinical LOD. However, since excess of information can become a disadvantage, we intend

⁷<http://rapgmsbgyim.github.io>

to provide filters, such as the causation frequency or the strength of the evidence.

Comparing to the previous implementation (Zamborlini et al., 2015b), the following explicit improvements are observed: (i) reasoning over action type hierarchy allows for detecting non-straightforward interactions (e.g. *DO administer NSAID* and *DO NOT administer Aspirin* are in contradiction because *Aspirin* specializes *NSAID*); (ii) new datasets (e.g. Sider) are added without need for writing specific rules for detecting external interactions; and (iii) causation frequency, belief strength (evidence level) and recommendation strength can be represented. Implicit improvement are: (i) a more maintainable and reusable implementation that will favor new features and datasets to be introduced in future work; (ii) the more reliable and/or relevant information can be select.

4 RELATED WORK

Formal languages proposed for representing clinical guidelines as “computer interpretable” ones (Annette ten Teije, 2008; Peleg, 2013) were not designed to handle the combination of multiple CIGs (Zamborlini et al., 2014a). An alternative solution is the development of alert systems that are independent of the CGs. Such Computerized Physician Order Entry systems (CPOE), are used to alert physicians about drug interactions (Ammenwerth et al., 2008). Despite the usefulness of these systems, a lot can be gained by tackling interactions between general recommendations on the outset, rather than employing drug-interaction alerts on the hospital floor.

We have investigated this issue in a series of work. In (Zamborlini et al., 2014a) we analyzed related work that addresses recommendation interactions in different levels. Our research focuses on what we called the CIG level, i.e. it accounts for the need to combine guidelines and handle interactions before applying them to a specific patient. This is the case when common co-occurring diseases are considered during guideline development, but could also be needed for uncommon co-occurring diseases in the practice setting. The related work (Jafarpour, 2013; López-Vallverdú et al., 2013; Wilk and Michalowski, 2014) has as their main drawback the need for defining specific rules for each interaction, e.g. *give aspirin & don't give aspirin* requires a specific rule and *give ibuprofen & don't give ibuprofen* requires another rule (a more detailed analysis in (Zamborlini et al., 2015b)). As a consequence, they do not provide ‘scalable’ support for combining guidelines, particularly more than two. Piovesan et al. (Pi-

Interaction Type	Interacting Recommendations/Beliefs	Source
AlternativeActions	['Avoid thrombi - by should Administer NSAID', 'Avoid thrombi - by should Administer Tramadol']	Internal
Contradiction	['Avoid gastrointestinal bleeding - by should-not Administer Aspirin', 'Avoid thrombi - by should Administer NSAID']	Internal
ExternalAlternativeAction	['Administer Epoprostenol always causes Low Blood Coagulation', 'Avoid thrombi - by should Administer NSAID', 'Avoid thrombi - by should Administer Tramadol']	DrugBank
ExternalIncompatibleActions	['Administer Aspirin incompatible with Administer Ibuprofen', 'Avoid thrombi - by should Administer NSAID', 'Reduce pain - by should Administer Ibuprofen']	DrugBank
ExternalIncompatibleEffects	['Administer Bendroflumethiazide always causes High Level Blood Sugar', 'Reduce blood pressure - by should Administer Thiazide', 'Reduce blood sugar level - by should Administer Insulin']	Sider
ExternalIncompatibleEffects	['Administer Ibuprofen always causes High Blood Pressure', 'Reduce blood pressure - by should Administer Thiazide', 'Reduce pain - by should Administer Ibuprofen']	Sider
ExternalAlternativeAction	['Administer Dipyridamole always causes Low Blood Coagulation', 'Avoid thrombi - by should Administer NSAID', 'Avoid thrombi - by should Administer Tramadol']	DrugBank
ExternalIncompatibleEffects	['Administer Ibuprofen always causes High Level Blood Sugar', 'Avoid thrombi - by should Administer NSAID', 'Reduce blood sugar level - by should Administer Insulin']	Sider

Figure 8: Case study on combining guidelines for OA+HT+DB.

ovesan et al., 2014) propose guideline-independent algorithms based on ontologies for detecting interactions, restricted to types “concordance” and “discordance”. The use of intentions associated to recommendations for detecting “intention interactions” is close to our approach on verifying transitions related to recommendations. To the best of our knowledge, none of the related work provides means to express negative norms, nor negative causation beliefs. Moreover, they do not explore action type hierarchies, nor reuse clinical knowledge available online in order to enrich the detection of interactions. However, they do address other aspects that we do not address yet, such as intentions, temporal aspects and qualitative transitions.

Our earlier work highlighted the importance of having the recommendations formally represented with a high level of detail. Explicit description of local constraints and impact of recommendations is considered an important source of information for increased reasoning capabilities and improved explanation of conflicts in (Bonacin et al., 2013). The model described in (Zamborlini et al., 2014a) introduced clinical recommendations as governing care actions that cause state transitions; an extended version of this model presented in (Zamborlini et al., 2014b) defines different ways in which recommendation can interact according to the referred actions and transitions. The implementation and evaluation of the model using Semantic Web languages was proposed by us in (Zamborlini et al., 2015b). We argued that the detection of interactions using *external* knowledge sources (in our case drug interactions modeled in the Linked Data version of DrugBank (Law et al., 2014)) can provide more precise information. A Web-based application for browsing the guideline interactions was made available online⁸. Extending this model to introduce the notion of causation beliefs (for evidence) and the

subsumption relations among actions was presented in (Zamborlini et al., 2015a). It was a first formal exercise with the goal of providing a systematic view on possible internal interactions among recommendations.

The emphasis on evidence means that care recommendations are ultimately grounded in domain knowledge (generalizations over facts). The evidence that underlies the recommendations is weighed depending on the quality, depth and breadth of the study: guidelines are part of a larger network of hypotheses, claims and pieces of evidence that span across multiple publications (de Waard et al., 2009). However, only few CIG languages offer means to link to evidence (Peleg, 2013), and they generally are targeted to very concrete and *procedural* guidelines, akin to medical protocols. In (Hoekstra et al., 2012), the authors describe a lightweight ontology that represents the relations between a guideline, its recommendations, and underlying evidence, as annotations on the guideline and evidence texts using a combination of the Open Annotation⁹ and PROV¹⁰ formats. Huang et al. (Huang et al., 2014) propose an even more lightweight semantic representation of evidence based clinical guidelines, but automatically extract it from guideline texts. It includes UMLS identifiers for medical terms appearing in the text and use proximity, and the types of terms to infer the type and strength of the evidence that underlies recommendations. The Nanopublication model (Mons et al., 2011) seems to be a natural fit to modeling the evidence that underlies guidelines. It represents a publication as three RDF graphs, that respectively capture an assertion (the finding or evidence), the provenance of the assertion (e.g. an experiment) and publication information about the nanopublication (when was the assertion published and by whom).

⁸See <http://guidelines.hoekstra.ops.few.vu.nl>

⁹See <http://www.openannotation.org>.

¹⁰See <http://www.w3.org/TR/prov-o>.

The work presented here combines the pragmatic approaches of (Hoekstra et al., 2012; Huang et al., 2014) and (Mons et al., 2011) in a model that takes the epistemological stance that the evidence underlying a recommendation expresses a *belief* that a care action causes a certain state transition. This strategy allows for using classical logic-based languages for handling inconsistent knowledge, such as conflicting findings published in different clinical studies.

5 DISCUSSION & CONCLUSION

The work reported on in this paper improves over our previous work by offering a more generic and scalable way to represent clinical guidelines and detecting interactions. This is done by adapting and extending both the conceptual model and the Semantic Web-based implementation. The TMR models and rules are made more generic so that they can be more easily extended to incorporate new features such as hierarchies of transition types and causal chains. Incorporating the epistemological nuance of beliefs in the Semantic Web representation, improves the ability to (i) handle knowledge from different sources and (ii) select the reliable ones; (iii) to allow different, incompatible beliefs about the same event to co-exist; and (iv) to provide reusable formal rules that are applicable regardless of specific regulations, guidelines or external sources. This has a favorable effect on reusability, maintainability, and scalability beyond the guidelines we currently covered.

We furthermore show the power of using the extensive domain knowledge available on the Semantic Web for enhancing the ability to automatically perform new tasks, such as suggesting alternative drugs. Our use of open standards and vocabularies, such as the nanopublication format, makes that the knowledge accumulated in our own models is shareable and reusable in a similar fashion. We implement inferencing using expressive SWI-Prolog rules that execute over RDF graphs. The adoption of SWI-Prolog was an improvement over the implementation in (Zamborlini et al., 2015b), as it gave us a single environment for expressing our inference rules benefiting understandability and maintainability. This, of course, at the cost of Semantic Web standards compliance for that specific part of our model. In (Zamborlini et al., 2015b) the limitations of OWL2 for detecting the interactions, forced the use of multiple knowledge representation languages. We had to resort to a combination of expressive OWL2 inferencing, Stardog SPARQL rules (a SWRL dialect) and custom SPARQL update queries to perform reasoning.

The experimental assessment shows that interactions can be automatically detected among three guidelines and enriched by knowledge from DrugBank and Sider, from each of which the relevant knowledge was imported as beliefs. Although the case study comprises only drug administration as action types, the approach is designed to address interactions among other types of interventions, such as surgeries and exercise therapy. More complex case studies will be addressed in future work. We faced some issues regarding the integration with these external knowledge sources, particularly on deciding which identity criteria we should use to map to the external datasets. For example, we could choose between PubChem ID, UMLS code, dbpedia and so on, where each choice would bring about different coverage and reliability.

Although this work is applied to clinical guidelines, its potentially of more general application, since both the model and the rules are defined independently of a particular domain. We plan to investigate the applicability of the models and rules to other domains such as disaster management. As ongoing work, we plan to address four limitations: (i) temporal validity for the assertions; (ii) quantification of beliefs and norms (i.e., frequency and strength); (iii) qualification of transitions (e.g. increasing or decreasing a property value); and (iv) considering goals and intentions;

ACKNOWLEDGEMENTS

The first author would like to thank colleagues from NEMO-UFES/Brazil for fruitful discussions about transitions, causation beliefs and regulations, and also prof. md. Saulo Bortolon for the nice discussions about medical domain; Jan Wielemaker and Wouter Beek from VU for helping with SWI-Prolog implementation; Wytze Vliestra (former VU) for fruitful discussions about the biomedical domain; and Paul Groth from Elsevier for fruitful discussions about the potential generality of the model and the use of nanopublications, This work was partially funded by the Dutch National Programme COMMIT.

REFERENCES

- Ammenwerth, E., Schnell-Inderst, P., Machan, C., and Siebert, U. (2008). The effect of electronic prescribing on medication errors and adverse drug events: a systematic review. *Journal of the American Medical Informatics Association*, 15(5):585–600.

- Annette ten Teije, Silvia Miksch, P. L., editor (2008). *Computer-based Medical Guidelines and Protocols: A Primer and Current Trends*, volume 139 of *Technology and Informatics*.
- Barnett, K., Mercer, S., Norbury, M., and Watt, G. (2012). Epidemiology of multimorbidity and implications for health care, research, and medical education: a cross-sectional study. *The Lancet*.
- Bonacin, R., Pruski, C., and Da Silveira, M. (2013). Architecture and services for formalising and evaluating care actions from computer-interpretable guidelines. *IJMEI International Journal of Medical Engineering and Informatics*, 5:253–268.
- de Waard, A., Shum, S. B., Carusi, A., Park, J., Samwald, M., and Sándor, Á. (2009). Hypotheses, evidence and relationships: The hyper approach for representing scientific knowledge claims. In *Proceedings of the 8th ISWC, Workshop on Semantic Web Applications in Scientific Discourse*, Berlin. Springer.
- Guizzardi, G., Wagner, G., de Almeida Falbo, R., Guizzardi, R. S. S., and Almeida, J. P. A. (2013). Towards Ontological Foundations for the Conceptual Modeling of Events. In *Conceptual Modeling, 32th International Conference, ER 2013*, pages 327–341, Hong-Kong. Springer Berlin Heidelberg.
- Hoekstra, R., de Waard, A., and Vdovjak, R. (2012). Annotating evidence based clinical guidelines - A lightweight ontology. In Paschke, A., Burger, A., Romano, P., Marshall, M. S., and Splendiani, A., editors, *Proceedings of the 5th International Workshop on Semantic Web Applications and Tools for Life Sciences, Paris, France, November 28-30, 2012*, volume 952 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Huang, Z., ten Teije, A., van Harmelen, F., and Ait-Mokhtar, S. (2014). Semantic Representation of Evidence-based Clinical Guidelines. In *6th International Workshop on Knowledge Representation for Health Care (KR4HC2014)*, volume 8903 of *LNCS*.
- Jafarpour, B. (2013). *Ontology Merging using Semantically-defined Merge Criteria and OWL Reasoning Services: Towards Execution-time Merging of Multiple Clinical Workflows to Handle Comorbidity*. PhD thesis, Dalhousie University.
- Law, V., Knox, C., Djoumbou, Y., Jewison, T., Guo, A. C., Liu, Y., MacIejewski, A., Arndt, D., Wilson, M., Neveu, V., Tang, A., Gabriel, G., Ly, C., Adamjee, S., Dame, Z. T., Han, B., Zhou, Y., and Wishart, D. S. (2014). DrugBank 4.0: Shedding new light on drug metabolism. *Nucleic Acids Research*, 42(D1):1091–1097. D1091-7, PubMed ID: 24203711.
- Lohr, K. N. (2003). Rating the strength of scientific evidence: relevance for quality improvement programs. *International Journal for Quality in Health Care*, 16(1):9–18.
- López-Vallverdú, J. A., Riaño, D., and Collado, A. (2013). Rule-based combination of comorbid treatments for chronic diseases applied to hypertension, diabetes mellitus and heart failure. In *LNCS*, volume 7738 *LNAI*, pages 30–41.
- Mons, B., van Haagen, H., Chichester, C., Hoen, P.-B., den Dunnen, J., van Ommen, G., van Mulligen, E., Singh, B., Hooft, R., Roos, M., Hammond, J., Kiesel, B., Giardine, B., Velterop, J., Groth, P., and Schultes, E. (2011). The value of data. *Nature Genetics*, 43(4):281–283.
- Peleg, M. (2013). Computer-interpretable clinical guidelines: a methodological review. *Journal of biomedical informatics*, 46(4):744–63.
- Piovesan, L., Molino, G., and Terenziani, P. (2014). An ontological knowledge and multiple abstraction level decision support system in healthcare. *Decision Analytics*, 1(1):8.
- Wilk, S. and Michalowski, M. (2014). Using First-Order Logic to Represent Clinical Practice Guidelines and to Mitigate Adverse Interactions. In *Knowledge Representation for Health-Care (KR4HC)*. *LNCS*, vol. 8903, Berlin Heidelberg. Springer.
- Zamborlini, V., da Silveira, M., Pruski, C., ten Teije, A., and van Harmelen, F. (2014a). Towards a Conceptual Model for Enhancing Reasoning about Clinical Guidelines: A case-study on Comorbidity. In *Knowledge Representation for Health-Care (KR4HC)*. *LNCS*, vol. 8903, Vienna, Austria. Springer Berlin Heidelberg.
- Zamborlini, V., da Silveira, M., Pruski, C., ten Teije, A., and van Harmelen, F. (2015a). Analyzing Recommendations Interactions in Clinical Guidelines: Impact of action type hierarchies and causation beliefs. In *Artificial Intelligence in Medicine (AIME)*. *LNCS*. Springer.
- Zamborlini, V., Hoekstra, R., da Silveira, M., Pruski, C., ten Teije, A., and van Harmelen, F. (2014b). A Conceptual Model for Detecting Interactions among Medical Recommendations in Clinical Guidelines. In *Knowledge Engineering and Knowledge Management (EKAW)*. *LNCS*, vol. 8876, pages 591–606. Springer.
- Zamborlini, V., Hoekstra, R., da Silveira, M., Pruski, C., ten Teije, A., and van Harmelen, F. (2015b). Inferring Recommendation Interactions in Clinical Guidelines: Case-studies on Multimorbidity. *Semantic Web Journal, Accepted, Open Access*.