Argumentation for Explainable Reasoning with Conflicting Medical Recommendations

Kristijonas Čyras, Brendan Delaney, Denys Prociuk, Francesca Toni
Imperial College London
London, UK

Martin Chapman, Jesús Domínguez, Vasa Curcin
King's College London
London, UK

Abstract

Designing a treatment path for a patient suffering from multiple conditions involves merging and applying multiple clinical guidelines and is recognised as a difficult task. This is especially relevant in the treatment of patients with multiple chronic diseases, such as chronic obstructive pulmonary disease, because of the high risk of any treatment change having potentially lethal exacerbations. Clinical guidelines are typically designed to assist a clinician in treating a single condition with no general method for integrating them. Additionally, guidelines for different conditions may contain mutually conflicting recommendations with certain actions potentially leading to adverse effects. Finally, individual patient preferences need to be respected when making decisions.

In this work we present a description of an integrated framework and a system to execute conflicting clinical guideline recommendations by taking into account patient specific information and preferences of various parties. Overall, our framework combines a patient's electronic health record data with clinical guideline representation to obtain personalised recommendations, uses computational argumentation techniques to resolve conflicts among recommendations while respecting preferences of various parties involved, if any, and yields conflict-free recommendations that are inspectable and explainable. The system implementing our framework will allow for continuous learning by taking feedback from the decision makers and integrating it within its pipeline.

1 Introduction

The Learning Health System (LHS), as defined by the US Institute of Medicine is a term that describes a system-wide approach to research and knowledge translation based on the exploitation of routinely collected data (McGinnis, 2010; McGinnis, Powers, and Grossmann, 2011). LHS is a part of a growing field of 'learning systems' where knowledge acquisition and process improvement become at least semi-automated tasks of the human-cyber-social infrastructure (Friedman et al., 2015). A number of projects and developments have made such a LHS in diagnostic and treatment decision support within reach, e.g. (Delaney et al., 2015) has built a prototype decision support system, integrated in UK primary care and shown a statistically significant improvement in diagnostic accuracy (Kostopoulou et al., 2017).

One significant component of any LHS is medical decision making and the possibilities of giving it automated support. In practice, medical decision making is often supported by clinical guidelines. These are lengthy documents summarising state-of-the-art knowledge about a medical condition and specifically its management. Clinical guidelines provide best practice recommendations to clinicians for taking care of patients under broad circumstances, mostly in the context of a single condition or a disease. Guidelines describe management of a generic patient, recommending multiple possible options for a clinician to opt for, depending on the specific circumstances. Such specific circumstances influencing the decision may amount not only to particularities regarding the disease in question, but also the presence of other diseases, i.e. the so called multimorbidity setting.

Multimorbidity (or comorbidity) is especially relevant in the context of chronic diseases, such as diabetes, asthma, chronic obstructive pulmonary disease (COPD), chronic kidney disease (CKD), to name a few. They are in some combination very often present for continuous periods of time, usually until death, especially in elderly patients. Furthermore, chronic diseases very often interact with each other in that managing one has positive and/or negatives effects on the others (Grace et al., 2013; Fraccaro et al., 2015). In particular, whereas certain clinical guideline recommendations are applicable for managing a given chronic disease in general, they are no longer such in the presence of other chronic diseases. What is more, due to the complexity of interactions, clinical guidelines that cover multimorbidities hardly ever exist. Thus, when managing multiple health conditions, different guidelines need to be considered.

Considering multiple guidelines very likely entails the existence of interacting recommendations: the recommendations may be inapplicable, may suggest incompatible actions or imply conflicting effects, may overlap, and so forth. A clinician may find it very difficult to follow the best practices should they stem from conflicting assumptions and/or lead to negative effects with respect to one or another condition. In order to facilitate the clinician's job, knowledge representation methods are useful in representing clinical guidelines and their interactions. However, whereas models for representing guidelines abound, see e.g. (Peleg, 2013) for an overview, few of them allow for capturing interactions among guideline recommendations. One state-of-the-art model that does have the latter feature is the Transition-based Medical Recommendation model (TMR) with the most recent exposition given by Zamborlini et al. (2017).

TMR allows for representing and merging multiple guidelines while taking into account their interactions. In the context of multimorbidities especially, TMR is very useful for identifying components and relations, such as clinical care actions, their positive and negative effects with respect to various conditions, as well as various measures of the quality of evidence and obligatoriness of recommendations. To capture interactions when merging guidelines, Zamborlini et al. (2017) advance a method to identify relationships among multiple recommendations, such as contradictions, side-effects, alternatives. The interactions are also accompanied with a measure of the degree of certainty that an interaction will happen. Therefore, TMR offers a detailed and comprehensive template for representing clinical guideline recommendations and their interactions. However, as TMR concerns generic recommendations, it does not afford a method for representing patient specific information. As a consequence of this, TMR does not possess a *reasoning* mechanism that would allow to determine the most applicable recommendations for a given patient.

Automated reasoning with clinical guideline representations and patient information, especially in the presence of guideline interactions, is an open problem in general (Peleg, 2013; Fraccaro et al., 2015). A further complication regarding reasoning with guideline recommendations and patient specific information is the need to take into account *preferences* of various parties involved – such as the patient, clinician and health care institution, see e.g. (Peleg, 2013; Wilk et al., 2017) for discussions. In this work we propose to apply an argumentation-based method for reasoning with clinical guidelines, patient information and various preferences.

Generally speaking, argumentation is a branch of knowledge representation and reasoning concerned with reasoning with partial and conflicting information in a way that aims to emulate human reasoning. In medical reasoning particularly, "argumentation is appealing as it allows for important conflicts to be highlighted and analysed and unimportant conflicts to be suppressed." (Atkinson et al., 2017) Structured argumentation formalisms—see e.g. (Besnard et al., 2014)—in particular provide ways to comprehensively represent information for reasoning medical knowledge via arguable elements and rules, see e.g. (Tolchinsky et al., 2006; Hunter and Williams, 2012). As such, argumentation formalisms are interpretable and naturally afford *explainable* reasoning methods. *Assumption-Based Argumentation with Preferences* (*ABA*$^+$) (Bondarenko et al., 1997; Čyras and Toni, 2016) is one established structured argumentation formalism that also deals with preference information. We propose to use ABA$^+$ for automating reasoning with conflicting guideline recommendations, patient specific information and preferences.

To enable this, we map TMR to the ABA$^+$ representation based on rules and arguable elements, called *assumptions*. This framework is instantiated using information extracted from computation representations of guidelines held in software that realises TMR. We also augment the representation in ABA$^+$ with patient specific conditions obtained from their electronic health record (EHR). This information allows us to construct arguments (as rule-based deductions) for actions based on recommendations. We also allow for the representation of preferences over assumptions, which influence how arguments and counterarguments interact (in argumentation jargon, *attack* each other). We then employ extension-based argumentation semantics to execute the reasoning and obtain the acceptable assumptions as well as arguments and conclusions. The reasoning outcomes are explainable through inspection of the explicitly given assumptions, rules, preferences as well as the resulting arguments and their relationships.

To illustrate our methodology, we focus on the interaction of conflicting recommendations. As an example, we consider an artificial case study of COPD, vetted by COPD experts. In the use case, a clinician deals with a patient that presents COPD and a mild Angina. The relevant clinical guidelines recommend several actions to take or avoid. We complement the recommendations with patient information from EHR and illustrate reasoning with and without preferences. We briefly discuss why the reasoning outcomes provided by ABA$^+$ are interpretable and explainable, and also discuss how the clinician can interact with a decision support system encompassing the ABA$^+$-driven reasoning engine.

This paper presents work in progress. Several parts of our LHS are in place, others are being researched and implemented. With this paper we aim to give a flavour of various parts and how they can come together to support an LHS.

The paper is structured thus. In Section 2 we give preliminaries about the TMR model, its implementation and integration with EHR data, as well as background on ABA$^+$. In Section 3 we advance a method for automated reasoning with guideline recommendations and patient information in ABA$^+$. We illustrate our approach with a COPD use case in Section 4. We discuss related work in Section 5 and conclude in Section 6.

2 Preliminaries

2.1 Transition-based Medical Recommendation (TMR) Model

In this section we review the Transition-based Medical Recommendation model (TMR) together with clinical guideline recommendation interaction description as a knowledge representation model. As in (Zamborlini et al., 2017), without loss of generality we assume that a set of guidelines is merged into a single guideline. We can thus assume that recommendations are delivered by the same larger guideline and avoid the need to refer to various guidelines.

Figure 1 depicts an instance of a graphical schema for representing recommendations in TMR. It consists of the following components.

- Unique *name* at the top of a rounded box. For instance, R_1, R_2. (We write R_k instead of Rk.)
  Henceforth, we refer to a recommendation by its name.
- Associated *action A* within the ellipse at the top. For instance, *Adm. NSAID*, *Adm. Aspirin*, where *Adm.* stands for Administer.
- *Deontic strength* indicated on the thick labelled arrow going out of the recommendation's name and into the action.

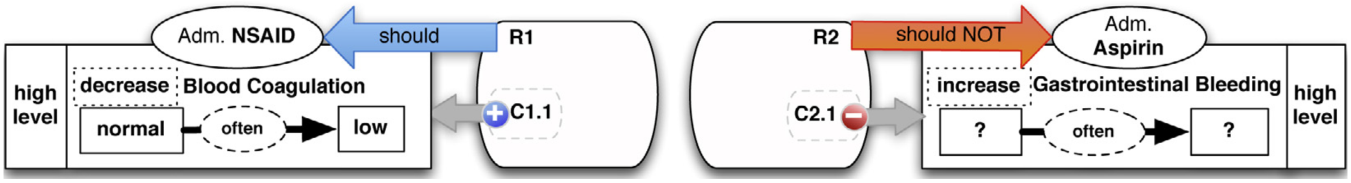

Figure 1: TMR representation schema instantiated with recommendations R_1 and R_2 (Zamborlini et al., 2017, p. 83, Figure 2).

For recommendation R, we denote its deontic strength by $\delta(R)$. It "reflects a degree of obligatoriness expected for that recommendation" (Zamborlini et al., 2017, p. 82). $\delta(R)$ takes values in $[-1, 1]$, being *positive* when $\geqslant 0$ and *negative* when < 0. If $\delta(R) \geqslant 0$, then R recommends to perform the action; else, if $\delta(R) < 0$, then R recommends to avoid the action.

As in (Zamborlini et al., 2017), to discretise $\delta(R)$ we may use the qualitative landmarks *must*, *should*, *may*, *should not* and *must not* corresponding to values 1, 0.5, 0, -0.5, -1, respectively. For instance, $\delta(R_1) = 0.5 =$ *should*, $\delta(R_2) = -0.5 =$ *should not*.

- *Properties* that the action affects, just below the action. For instance, *Blood Coag.* and *Gastro. Bleeding*. (We abbreviate words: e.g. *Gastro. Bleeding* abbreviates *Gastrointestinal Bleeding*.)

  In general, an action can affect more than one property P.
- *Effects* of the actions within the dashed rectangles to the left of the properties. For instance, *decrease* and *increase*. An action A has one effect E on the property P it affects. Effects may have determinate *initial* and *final* values, within the rectangular boxes below the property in question, the black arrow coming out of the initial value (box) and leading into the final value (box). Otherwise, ? represents an *indeterminate* value.

  For instance, action *Adm. NSAID* affects *Blood Coag.* by decreasing it from the initial value *normal* to the final value *low*. On the other hand, *Adm. Aspirin* increases *Gastro. Bleeding* with indeterminate values.

  In this paper we *will not* make use of, but mention for completeness, two quantitative values associated with an effect. One is the *causation probability* within the dashed ellipse below the property. It represents the likelihood of the action bringing the effect about. For instance, *often*.

  The other one is the *belief strength* boxed to the left- or right-most side. It represents the level of evidence regarding bringing the effect about. For instance, *normal level*.
- *Contributions* of the recommendation to the overall goals in the context of a guideline indicated below the recommendation name within a transparent dashed rounded box. For instance, $+C1.1$, $-C2.1$.

  In general, recommendation R can have more than one contribution. Each contribution carries an *identifier*, e.g. $C1.1$, $C2.1$, and is *valued* in $[-1, 1]$, depending on how important it is to achieve or avoid the corresponding effect. The value is discretised with signs: $+$, $-$ and the absence of a sign represent, respectively, values greater than, less than and equal to 0.

  As in (Zamborlini et al., 2017), the overall goal is to always abstract "patient well-being". However, for a given patient, the clinician may, and in general will, have intermediate goals, such as to *decrease Blood Coag.* In this paper, we are not specifically concerned with intermediate goals and take them to be implicitly given by effects that actions bring about.

We will use $\mathbb{R}$ to denote a fixed but otherwise arbitrary set of recommendations.

Observe that an instance of TMR concerns a generic patient. In order to apply recommendations, one needs to consider specific patient *conditions*. Such conditions pertain to properties and the initial values of the effects that actions have on properties. For instance, a patient can have *normal Blood Coag.* or *Gastro. Bleeding*. When using ABA$^+$ to reason with guidelines, patient conditions will come as information additional to TMR instances.

Using TMR, Zamborlini et al. (2017) identify *interactions* among recommendations. Intuitively, interactions record the relationships between different recommendations. Several types of interactions are possible, namely contradiction, repetition, alternative, side-effect, repairable and safety. For instance, a *contradiction* interaction arises between two recommendations if one states that the action suggested by the other should be avoided. A *side-effect* interaction arises when the action of one recommendation causes a secondary effect which is opposite to the effect of the action of the other recommendation. For example, Ibuprofen may *increase Blood Press.*, which is aimed to be decreased by another medication. Not all interactions concern conflicting relationships, though. For instance, *repetition* indicates that two recommendations suggest (roughly) the same course of action, e.g. *Adm. NSAID* and *Adm. Aspirin*.

The implementation of TMR used will allow for all such interactions. However, for the purpose of illustrating reasoning with recommendations using argumentation, we focus on the contradiction interaction in this paper, because it relates recommendations in direct conflict that can be naturally resolved by means of argumentation.

Formally, interactions can be represented as triples (R, R', μ) with recommendations R and R', and the interaction's *modal strength* μ, which reflects the conclusiveness of the interaction. The interaction's modal strength can take two values, denoted by $\Box$ and $\Diamond$, where $\Box$ means "the interaction will certainly occur if the related recommendations are prescribed" (Zamborlini et al., 2017) and $\Diamond$ means "the interaction is uncertain to happen". We assume that the interactions and their modal strengths are given along with the instances of the TMR model. We will use $\mathbb{I}$ to denote the set of all (contradiction) interactions given $\mathbb{R}$.

Note well that $\mathbb{R}$ and $\mathbb{I}$ amount only to representation of guidelines, but not reasoning with them. In particular, it is a patient-agnostic representation, while the reasoning happens with patient-specific information. The following example illustrates recommendations and their interactions.

Example 2.1. The two recommendations R_1 and R_2 as in Figure 1 can be considered in (contradiction) interaction, because they recommend opposite actions.[1] So let $\mathbb{R} = \{R_1, R_2\}$ and assume $\mathbb{I} = \{(R_1, R_2, \Box)\}$. Intuitively, for a generic patient, NSAID—e.g. Aspirin—should be administered. If, however, the patient exhibits *Gastro. Bleeding*, then R_1 and R_2 are in conflict and there are arguments for both administering and not administering Aspirin.

To resolve the conflict in this case, one could administer a different NSAID, such as Ibuprofen. However, in more complicated situations such alternatives may not be readily available, whence certain actions should not be taken (i.e. certain recommendations cannot be followed).

2.2 Guideline and EHR Data

An implementation of TMR (see for instance `guidelines2.eculture.labs.vu.nl/swish/p/datasetMaintenance.swinb`) allows for the computational representation of clinical guidelines using standards such as the Resource Description Framework. We are thus able to represent guidelines in a manner that makes them amenable to the automatic instantiation of ABA$^+$ frameworks for reasoning with guideline recommendations.

Similarly, the EHR data required to reason with guidelines in light of patient-specific information can be extracted automatically via pieces of middleware running within GPs' practices. These pieces of middleware are designed to communicate with the Application Programming Interface of a locally installed EHR system, and communicate patient record information, at the discretion of the practitioner, to us for use. Patient information can then be modelled as part of a given TMR implementation, such as in the form of a set of additional external rules, prior to being used in the reasoning process, or delivered to the reasoning engine separately.

This flow of data creates a *decision support pipeline*, in which potentially conflicting guideline data, along with EHR data, is passed to a reasoning engine that returns non-conflicting recommendations for use by the system.

2.3 Assumption-Based Argumentation with Preferences (ABA$^+$)

Argumentation is a branch of the field of Artificial Intelligence concerned with reasoning with partial and conflicting information. For example, medical information is often partial, because it may be infeasible or simply unreasonable to record all the possibly relevant patient information; medical information can also be conflicting, as in the case of conflicting clinical guideline recommendations. As such, argumentation lends itself to be applied for reasoning purposes in the context of guidelines and patient specific information.

An important feature of many argumentation formalisms is that they are inherently interpretable and afford explainable reasoning. What this amounts to is the construction of arguments and counterarguments for explicit claims, based on explicit assumptions, using rules expressing application specific reasoning patterns. In addition to providing inspectable arguments and counterarguments for or against claims that may encode beliefs, decisions etc., argumentation allows for questioning the assumptions underlying the arguments and for a continuous addition of new assumptions (and thus arguments). With this, one can explain why e.g. a particular decision was taken, i.e. what were the arguments for and against it, and how can one further question and/or support that decision.

In this work we use a well-established and broadly-studied argumentation formalism, called ABA (Bondarenko et al., 1997), and its extension ABA$^+$ (Čyras and Toni, 2016; Bao, Čyras, and Toni, 2017), because it has all the features discussed above. We provide the background for ABA$^+$ following (Bondarenko et al., 1997; Čyras and Toni, 2016).

An *ABA$^+$ framework* is a tuple $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ }, \leqslant)$, where:

- $(\mathcal{L}, \mathcal{R})$ is a deductive system with $\mathcal{L}$ a language and $\mathcal{R}$ a set of rules of the form $\varphi_0 \leftarrow \varphi_1, \ldots, \varphi_m$ with $m \geqslant 1$, or of the form $\varphi_0 \leftarrow \top$, where $\varphi_i \in \mathcal{L}$ for $i \in \{0, \ldots, m\}$ and $\top \notin \mathcal{L}$; φ_0 is the *head* and $\varphi_1, \ldots, \varphi_m$ the *body* of the rule; $\varphi_0 \leftarrow \top$ is said to have an empty body and called a *fact*;
- $\mathcal{A} \subseteq \mathcal{L}$ is a non-empty set of *assumptions*;
- $\bar{\ }: \mathcal{A} \to \mathcal{L}$ is a total map: for $\alpha \in \mathcal{A}$, $\overline{\alpha}$ is referred to as the *contrary* of α;
- $\leqslant$ is a preorder (i.e. reflexive and transitive order) on $\mathcal{A}$, called a *preference relation*.

As usual, the strict (asymmetric) counterpart $<$ of $\leqslant$ is given by $\alpha < \beta$ iff $\alpha \leqslant \beta$ and $\beta \nleqslant \alpha$, for any α and β. (We assume this for all preorders in this paper.) For assumptions $\alpha, \beta \in \mathcal{A}$, $\alpha \leqslant \beta$ means that β is at least as preferred as α, and $\alpha < \beta$ means that α is strictly less preferred than β.

Throughout the paper, we assume as given a fixed but otherwise arbitrary ABA$^+$ framework $\mathcal{F} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ }, \leqslant)$, unless specified otherwise. If the preference relation $\leqslant$ in $\mathcal{F} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ }, \leqslant)$ is empty or unspecified, i.e. there are no preferences, then we may refer to $\mathcal{F}$ as an *ABA framework* Bondarenko et al. (1997) and denote it $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ })$.

Assumptions in ABA$^+$ represent arguable information. For instance, assumptions can represent the applicability of, or an agent's willingness to follow, a recommendation. In such a case, preferences in ABA$^+$ can represent the relative degrees obligatoriness, or willingness to follow, the recommendations. We will exemplify various ABA$^+$ components in Section 3. We next give notions of arguments (as deduction trees) and attacks in ABA$^+$.

An *argument for $\varphi \in \mathcal{L}$ supported by $A \subseteq \mathcal{A}$ and $R \subseteq \mathcal{R}$*, denoted $A \vdash^R \varphi$, is a finite tree with: the root labelled by φ; leaves labelled by $\top$ or assumptions, with A being the set of all such assumptions; the children of non-leaves ψ labelled

[1] Note well that a hierarchy of actions is assumed (Zamborlini et al., 2017, p. 79) to obtain interactions. For instance, the action to administer NSAID subsumes both actions to administer Aspirin and Ibuprofen. Such a hierarchy can be accessed via queries to a populated implementation of TMR (See Section 2.2).

by the elements of the body of some ψ-headed rule in $\mathcal{R}$, with R being the set of all such rules. $A \vdash \varphi$ is a shorthand for an argument $A \vdash^R \varphi$ with some $R \subseteq \mathcal{R}$.

For $A, B \subseteq \mathcal{A}$, A *<-attacks* B, denoted $A \rightsquigarrow_< B$, iff:

a) either there is an argument $A' \vdash \overline{\beta}$, for some $\beta \in B$, supported by $A' \subseteq A$, and $\nexists \alpha' \in A'$ with $\alpha' < \beta$;
b) or there is an argument $B' \vdash \overline{\alpha}$, for some $\alpha \in A$, supported by $B' \subseteq B$, and $\exists \beta' \in B'$ with $\beta' < \alpha$.

The intuition here is that A <-attacks B if a) either A argues contra something in B by means of no inferior elements (*normal attack*), b) or B argues contra something in A but with at least one inferior element (*reverse attack*).

If A does not <-attack B, we may write $A \not\rightsquigarrow_< B$. For ABA frameworks $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ we often drop the subscript/prefix $<$ and say e.g. *attacks*, written $\rightsquigarrow$. Note that without preferences, an attack from one set of assumptions to another boils down to the former set deducing the contrary of some assumption in the latter set.

In the setting of guideline recommendations, sets of assumptions will represent sets of recommendations, and they will induce arguments for or against following recommendations. The attacks among sets of assumptions (recommendations) will arise due to existence of interactions. The preferences may come from e.g. the patient or the clinician.

The reasoning in ABA$^+$ is realised through the *semantics*. Intuitively, a semantics gives conditions that a set of assumptions needs to satisfy in order to be 'acceptable', or 'good'. The conclusions derived from acceptable assumptions represent a coherent set of beliefs, decisions to make, actions to take, etc., depending on the problem formulation, and are thus deemed as the reasoning outcomes. We next give notions used to define ABA$^+$ semantics. Let $A \subseteq \mathcal{A}$.

- The *conclusions of* A is the set $Cn(A) = \{\varphi \in \mathcal{L} : \exists A' \vdash \varphi, A' \subseteq A\}$ of sentences concluded by (arguments supported by subsets of) A.

We next give three basic requirements for sets of assumptions to be 'good', or collectively acceptable. The first one says such a set should include all assumptions it makes.

1. We say A is *closed* if $A = Cn(A) \cap \mathcal{A}$, i.e. A contains all assumptions it concludes.

We say $\mathcal{F}$ is *flat* if every $A \subseteq \mathcal{A}$ is closed. We assume ABA$^+$ frameworks to be flat, unless specified otherwise.

The second one expresses that to be acceptable, a set should not be conflicting, i.e. not to <-attack itself.

2. A is *<-conflict-free* if $A \not\rightsquigarrow_< A$.

The third says that a 'good' set should defend against counterarguments; we first define the notion of defense.

- A *<-defends* $A' \subseteq \mathcal{A}$ if for all $B \subseteq \mathcal{A}$ with $B \rightsquigarrow_< A'$ it holds that $A \rightsquigarrow_< B$.

So finally,

3. A is *<-admissible* if it is <-conflict-free and <-defends itself.

These are arguably three 'minimal' requirements for accepting a given set of assumptions (as well as arguments based on them). Note that they are quite weak, because, for instance, the empty set of assumptions is always <-admissible. However, not much can in general be concluded from $\emptyset$. Thus, ABA$^+$ *semantics* impose additional requirements for acceptance of assumptions and the associated conclusions. In this paper we use one particular such semantics which says that a 'good' set of assumptions should be as large as possible, as follows.

- A set $E \subseteq \mathcal{A}$ of assumptions is a *<-preferred extension* of $\mathcal{F} = (\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{}, \leqslant)$ if E is $\subseteq$-maximally <-admissible.

In other words, with <-preferred extensions we are aiming to conclude as much as we can without contradicting ourselves, whilst being able to defend ourselves.

For ABA frameworks we often drop the prefix $<$ for the notions above.

3 Mapping TMR and EHR to ABA$^+$

In this section we discuss a mapping from TMR to ABA$^+$, augmenting the guideline recommendations and their interactions with patient specific information based on EHR.

For the purpose of mapping TMR instances to ABA$^+$, we assume a simplified TMR whose instances are recommendations given as tuples $(R, A, \delta(R), \mathcal{P}, \mathcal{E}, \mathcal{V}, \mathcal{C})$ with

(i) name R,
(ii) action A,
(iii) deontic strength $\delta(R)$,
(iv) properties $\mathcal{P} = \langle P^1, \ldots, P^n \rangle$, for $n \geqslant 1$,
(v) effects $\mathcal{E} = \langle E^1, \ldots, E^n \rangle$,
(vi) initial values $\mathcal{V} = \langle v^1, \ldots, v^n \rangle$ of effects on properties,
(vii) contribution values $\mathcal{C} = \langle c^1, \ldots, c^n \rangle$.

We identify any such recommendation with its name R.

We next describe a mapping from TMR to ABA$^+$. We omit the cumbersome formal details in the interest of space. We will exemplify the mapping in Section 4.

Given a recommendation $(R, A, \delta(R), \mathcal{P}, \mathcal{E}, \mathcal{V}, \mathcal{C})$, we construct the following:

a) an assumption $R \in \mathcal{A}$ representing the possible applicability of the recommendation;
b) a rule $A \leftarrow R \in \mathcal{R}$ representing that action A is recommended by R;
c) for each property P^i and its corresponding effect E^i, a rule $E^i P^i \leftarrow A \in \mathcal{R}$ representing that action A brings about effect E^i to property P^i.

We use the additional components of recommendations to model interactions. Specifically, suppose recommendations $(R_1, A_1, \delta(R_1), \mathcal{P}_1, \mathcal{E}_1, \mathcal{V}_1, \mathcal{C}_1)$ and $(R_2, A_2, \delta(R_2), \mathcal{P}_2, \mathcal{E}_2, \mathcal{V}_2, \mathcal{C}_2)$ are in contradiction, with actions A_1 and A_2 recommended positively ($\delta(R_1) > 0$) and negatively ($\delta(R_2) < 0$), respectively. That means R_2 can be argued against on the basis of R_1 and the presence of the interaction. On the other hand, R_1 can be similarly argued against on the basis of R_2 and the presence of the interaction, but only if a given patient presents some condition affected by A_2 that contributes negatively to the patient's well-being.

Thus, given $(R_1, R_2, \mu) \in \mathbb{I}$, we construct the following:

d) $\overline{R_2} \leftarrow R_1, int_{1,2}$;
e) $\overline{R_1} \leftarrow R_2, int_{1,2}, vP$,

where $P \in \mathcal{P}_2$ is a property with initial value $v \in \mathcal{V}_2$ and contribution $- = c \in \mathcal{C}_2$. (When the initial value v of P is indeterminate ?, we use only P in the body of the rule.) Here, $int_{1,2} \in \mathcal{L}$ represents $(R_1, R_2, \mu) \in \mathbb{I}$. The rule in d) says R_2 should not be followed if (i) R_1 is followed, and (ii) R_1 and

R_2 are in contradiction. The rule in e) says R_1 should not be followed if (i) R_2 is followed, (ii) R_1 and R_2 are in contradiction, *and also* (iii) the condition vP is present.

The interaction's modal strength determines whether the interaction can be argued about or not:

f) given μ,
   1. if $\mu = \Box$, let $int \leftarrow \top \in \mathcal{R}$;
   2. if $\mu = \Diamond$, let $int \in \mathcal{A}$.

The rule $int \leftarrow \top \in \mathcal{R}$ represents that the interaction is sure to happen, i.e. it is a fact, and so there is no way to disagree with it. However, as an assumption, $int \in \mathcal{A}$ represents that the interaction is not certain to happen and so can be argued against by putting forward arguments for the contrary $\overline{int}$.

Now, the patient specific conditions can be similarly represented as either facts or assumptions:

g) given a patient condition *cond*,
   1. either let $cond \leftarrow \top \in \mathcal{R}$;
   2. or let $cond \in \mathcal{A}$.

Whether the conditions can be argued about or not depends on the context. For instance, it may be debated whether a patient is taking certain medications (confirmation of which is part of standard hospital procedures), but it may be certain that a patient has mild Angina.

Of particular interest are those conditions that appear within recommendations as properties affected by actions. Specifically, if some condition as property P (and possibly value v) matches that in a recommendation, then the addition of P (or vP) as either a fact or an assumption may trigger a rule concerning interactions of recommendations, such as e) above. In this way a patient's EHR can meaningfully augment the TMR model when represented in ABA$^+$.

Given the assumptions and rules constructed from recommendations in $\mathbb{R}$ and interactions in $\mathbb{I}$ as per points a)–g) above, we can define an ABA framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ with contraries on assumptions $\alpha \in \mathcal{A}$ being new symbols $\overline{\alpha}$ and the language $\mathcal{L}$ given given by the symbols appearing in $\mathcal{A}$, $\mathcal{R}$ and $\{\overline{\alpha} : \alpha \in \mathcal{A}\}$. In $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ we can construct arguments and counterarguments for actions based on (the possibility of following) recommendations and patient specific conditions. The semantics (of e.g. preferred extensions, see Section 2.3) then allow to determine sets (i.e. extensions) of collectively non-conflicting recommendations for a given patient. The conclusions of such extensions then yield, in addition to the recommendations to be followed, the actions to be taken as well as their consequences in terms of effects on the patient's conditions.

In addition to recommendations and patient EHR information, we may also have *preferences* over e.g. courses of action, of one or the other party involved. For instance: the patient may prefer one medicine over another, according to what they are used to; or the clinician may prioritise one course of action over another, based on their professional experience; or else, the hospital may have preferences over treatment methods, judging by the information on their success in the local geographical region.

Preferences can be naturally incorporated in ABA$^+$ by extending the ABA framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ with a preference relation $\leqslant$ to obtain an ABA$^+$ framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{}, \leqslant)$. For example, if action A_1 suggested by recommendation R_1 is preferred over action A_2 suggested by recommendation R_2, then the preference $R_2 < R_1$ can be added. Such preferences then possibly affect the reasoning outcomes in that the extensions (and the associated conclusions) obtained respect the preferences specified.

Last but not least, the reasoning process in ABA$^+$ instantiated with TMR, EHR and preferences is fully transparent and explainable. Indeed, the assumptions and rules on which the arguments are based are clearly stated, the attack relationship among arguments is constructively defined based on the explicitly given assumptions and preferences, and the semantics comprehensively express reasonable requirement for argument/assumption acceptance.

We illustrate the mapping to and reasoning in ABA$^+$ with a use case in the next section.

4 COPD Use Case

The COPD use case developed sets up the boundaries and scope of the problem that is mapped and resolved throughout the rest of the model. The use case is established as existing within the context of a secondary health-care system and creates an artificial scenario of a patient that would present themselves within the health-care system with symptoms typical to COPD, and illustrate how they would be managed within the health-care system with regards to following official guideline recommendations.

4.1 Stable COPD with Conflicting Recommendations

The patient presents to a primary health-care setting complaining of breathing difficulties and increased fatigue during exercise, and following standard diagnostic procedures not covered by the model, is diagnosed with mild stable COPD. As for the treatment choice, a decision for medication is made based on the patients vitals as some of the main deciding factors in the form of: (i) Blood gas levels; (ii) Spirometry results; (iii) Age; (iv) Current lifestyle habits; (v) Existing comorbidities. For the purpose of illustrating the argumentation component of the model, the patient is assumed to have a pre-existing condition in the form of *mild Angina*, which was diagnosed at an earlier point in time.

Following the NICE COPD management guidelines (NICE, 2010), the patient is on schedule to be prescribed a short course of Short-acting Beta Agonists (*SABA*). But as outlined in the GOLD COPD management guideline (GOLD, 2017), which is used in conjunction with the NICE guideline in primary health-care in the UK, a patient presenting with angina should not be prescribed a standard SABA inhaler as it may lead to further exacerbation of their cardiovascular symptoms and progress towards heart failure. The guideline instead suggests the prescription of a reduced nebulised dose of the SABA medication, which while having the therapeutic effect intended on relieving symptoms of mild stable COPD, should not cause as much an irritation to the patients cardiovascular system.

Similarly, for patients who present with a mild form of COPD, the NICE guideline recommends that the patient undertake regular exercise to boost the functioning of their cardiovascular system. It is easy to imagine a situation where this would not be applicable, for example if the patient is suffering from joint pain, or peripheral artery disease. In such a case, the expected prescription would not only be unhelpful, it would in fact be damaging to the health of the patient. A contraindication to the aforementioned exercise can be found within the clinical guidelines for the specific comorbidities, highlighting the near unlimited potential complexity of a clinical course.

As such, the stable COPD with mild angina use case serves as a simple illustration of the concept that can thereupon be further expanded to encompass multimorbidities and conflicts in medication, and outpatient management.

4.2 Use Case in ABA$^+$

For illustrating how ABA$^+$ deals with interacting guideline recommendations, we take two recommendations from the COPD use case, namely administering SABA and not administering standard SABA. The two recommendations are in contradiction, because *standard SABA* is subsumed by *SABA*. Thus, $\mathbb{R}$ consists of the following recommendations.

1. $(R_1, A_1, \delta(R_1), \mathcal{P}_1, \mathcal{E}_1, \mathcal{V}_1, \mathcal{C}_1)$ with: (i) name R_1; (ii) $A_1 = \textit{SABA}$; (iii) $\delta(R_1) = \textit{must}$; (iv) $\mathcal{P}_1 = \langle \textit{Lung muscles}, \textit{Airways} \rangle$; (v) $\mathcal{E}_1 = \langle \textit{relaxes}, \textit{dilates} \rangle$; (vi) $\mathcal{V}_1 = \langle ?, ? \rangle$; (vii) $\mathcal{C}_1 = \langle +, + \rangle$;
2. $(R_2, A_2, \delta(R_2), \mathcal{P}_2, \mathcal{E}_2, \mathcal{V}_2, \mathcal{C}_2)$ with: (i) name R_2; (ii) *standard SABA*; (iii) $\delta(R_2) = \textit{must not}$; (iv) $\mathcal{P}_2 = \langle \textit{Angina} \rangle$; (v) $\mathcal{E}_2 = \langle \textit{increase} \rangle$; (vi) $\mathcal{V}_2 = \langle \textit{mild} \rangle$; (vii) $\mathcal{C}_2 = \langle - \rangle$.

We assume that the interaction between R_1 and R_2 is certain, so that $\mathbb{I} = \{(R_1, R_2, \Box)\}$. Importantly, $\mathbb{R}$ and $\mathbb{I}$ yield the assumptions $R_1, R_2 \in \mathcal{A}$ and the following rules in $\mathcal{R}$:

- $\overline{R_2} \leftarrow R_1, \textit{int}_{1,2}$;
- $\overline{R_1} \leftarrow R_2, \textit{int}_{1,2}, \textit{mild Angina}$;
- $\textit{int}_{1,2} \leftarrow \top$.

(For simplicity, we omit to specify the rules regarding the actions as well as their effects on properties.)

Given that the patient has a mild Angina, we have

- $\textit{mild Angina} \leftarrow \top$

in $\mathcal{R}$, representing the patient specific condition.

In the resulting ABA framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$ we find arguments $\{R_1\} \vdash \overline{R_2}$ and $\{R_2\} \vdash \overline{R_1}$, so that $\{R_1\}$ and $\{R_2\}$ attack each other. The two sets are thus preferred extensions of $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{})$, concluding respectively administering *SABA* and not administering *standard SABA*. As one of the conclusions is not to take any action, one can either employ preferences, or pass the information back to the TMR implementation to refine recommendations, if possible.

Regarding preferences, the clinician could insist that not worsening *Angina* takes priority over addressing COPD by way of administering *standard SABA*. Thus, the preference $R_1 < R_2$ could be added to obtain the ABA$^+$ framework $(\mathcal{L}, \mathcal{R}, \mathcal{A}, \overline{}, \leqslant)$. There, $\{R_2\} \leadsto_< \{R_1\}$, but $\{R_1\} \not\leadsto_< \{R_2\}$, so that R_2 forms a unique $<$-preferred extension and no action is recommended.

Following this, or otherwise prior to employing preferences, one can look for a refinement of the generic recommendation R_1. And indeed, recommendation R_3 can be found that is like R_1, but suggests administering *nebulised SABA* instead. Adding R_3 to the existing ABA/ABA$^+$ framework, or otherwise constructing a new one with R_3 replacing R_1, leads to obtaining $<$-preferred extensions which conclude administering *nebulised SABA*. This is in agreement with what should actually be done.

Note well that reasoning with conflicting information (as well as preferences) and yielding non-conflicting conclusions is not the only thing allowed by argumentation. In addition, argumentation affords means to inspect and explain the reasoning. In particular, the extensions obtained, as well as arguments for specific claims and/or based on specific assumptions, can be presented to the clinician or more generally a user of the LHS. This way the user can interact and provide feedback to the system so that it evolves and yields better reasoning outcomes in the future. We leave the description and implementation of feedback integration within the system for future work.

5 Related Work

Argumentation has already been successfully applied in health-care, see e.g. (Longo, 2016; Atkinson et al., 2017) for overviews. Different works can be distinguished by the components of the argumentative reasoning process they use.

There are several works that use both argument construction and argumentation semantics for reasoning with medical knowledge, as we do in this paper.

For instance, Hunter and Williams (2012) use a structured argumentation formalism and employ preferences to reason with conflicting medical knowledge. In their work, evidence from clinical trials is manually extracted from guidelines and synthesised to form arguments for, and counterarguments against, treatment superiority. Based on treatment outcome indicators and the importance of evidence, user-specified preferences over arguments are formed. Semantics of grounded (Dung, 1995) and preferred extensions are used to identify the acceptable arguments and thus the superior treatments. We, in contrast, focus on resolving conflicts among guideline recommendations when managing multimorbidities, rather than determining treatment superiority based on clinical trials. We also aim our methodology to yield explainable decision support.

The recent CONSULT project (Kokciyan et al., 2018) applies argumentation to reason with guidelines and patient preferences for managing post-stroke patients. Kokciyan et al. (2018) manually represent guidelines in first-order logic (FOL) and use argument schemes (Walton, 1996), preferences and argumentation semantics to resolve inconsistencies among recommendations. We instead build argumentation on the well-established TMR model and offer explainable decision making. We leave formal comparison with (Kokciyan et al., 2018) for future work.

Other works incorporating argumentation and preferences focus on helping clinicians to construct and evaluate arguments for and against decisions. As such, they do not automatically populate their argumentation frameworks with guideline knowledge or EHR data, and do concern reasoning with clinical guideline recommendations, but are nevertheless related to our work due to the use of argumentation semantics for reasoning.

For instance, Tolchinsky et al. (2006) use argumentation, its semantics and preferences in a multi-agent deliberation about organ transplantation. There, expert clinicians use argumentation schemes to construct arguments and attacks concerning viability of transplantation. A mediator agent then evaluates the arguments so as to determine their strength. The mediator agents does this by using as preferences over arguments the knowledge from clinical guidelines, as well as knowledge about past transplantations and the interacting agents themselves. Somewhat similar in spirit, the system ArgMed (Qassas et al., 2016) allows to document and turn clinicians' discussions into argumentation frameworks using argumentation schemes. After that, preferred semantics is used to determine the acceptable arguments and hence the best claims made by the clinicians.

In some works that use argumentation components to model information, argumentation semantics are not used to execute the reasoning itself. For instance, in one of the earliest related works, Fox et al. (2006) enable agents to exchange arguments in order to automate medical reasoning, albeit not with guideline recommendations. Arguments are assigned strength and the strengths can be aggregated using e.g. probabilistic or decision making approaches to determine the strongest arguments.

An argument aggregation mechanism for reasoning with guidelines is used in (Grando, Glasspool, and Boxwala, 2012). There, templates for generating arguments are based on argumentation schemes. Arguments roughly correspond to statements in clinical guidelines: an argument consists of assumptions, claim, polarity (for or against claim), confidence (representing, for instance, quality of the evidence or the likelihood of an outcome) and precondition (whether the argument is applicable). A unique goal needs to be specified when aggregating argument confidence metrics to reason about the strength of the arguments that enable one to achieve the goal in question. Aside from the use of argumentation semantics instead of argument aggregation, a few points make our work different: i) we focus on reasoning with conflicting recommendations from multiple guidelines, whereas Grando, Glasspool, and Boxwala (2012) are executing recommendations of a single guideline; ii) also, reasoning in ABA^+ is assumption-, rather than goal-, driven.

As for non-argumentative approaches to reasoning with interacting guidelines, Wilk et al. (2017) propose a framework for mitigating concurrent execution of clinical guidelines. They also deal with patient specific conditions and patient preferences. There, recommendations are represented as actionable graphs. Wilk et al. (2017) map those into FOL rules, and introduce patient conditions and preferences via FOL revision operators. Guideline mitigation then amounts to applying revision operators to FOL rules representing the recommendations, so as to account for patient specific conditions and preferences. Finally, reasoning is done by finding models of the resulting FOL theory.

Our work is different in terms of both reasoning and representation. Regarding representation, as indicated by Wilk et al. (2017), the TMR model is in some aspects richer than the mitigation specific FOL (however, TMR does not have a temporal component which is present in (Wilk et al., 2017)). Regarding reasoning mechanisms, model finding in FOL is in general an undecidable problem, as opposed to finding preferred extensions in ABA^+ frameworks. We also believe argumentation-based reasoning to be more transparent, as one can inspect the arguments, attacks among them and their interplay with preferences, in contrast to interpreting workings and results of a FOL theorem prover utilised by Wilk et al. (2017). It would be interesting though to integrate the temporal aspect into our implementation of the TMR model and within ABA^+. We leave this for future work.

Other approaches to reasoning with guidelines and temporal as well as clinical constraints exist (Peleg, 2013), mainly using task network models, see e.g. Leonardi et al. (2012); Shalom, Shahar, and Lunenfeld (2016). However, they deal with single rather than multiple guidelines and are thus not specialised to handle conflicts, as opposed to our approach.

6 Conclusions and Future Work

We described work in progress towards a Decision Support System that will use Transition-based Medical Recommendation model (TMR) and its integration with electronic health record (EHR) data to facilitate automated execution of interacting clinical guidelines by taking into account patient's individual medical history and preferences of various parties involved. In particular, we proposed the structured argumentation formalism ABA^+ for automated reasoning with conflicting clinical guideline recommendations as well as patient information and preferences. We also discussed how ABA^+ yields interpretable and explainable medical decisions for execution of guideline recommendations.

Future work on implementations of TMR involves constructing standard interfaces that can be queried ad-hoc by other systems such as argumentation frameworks. Similar efforts to increase the interoperability of EHR data include the examination of techniques to standardise independent vendor formats to a single target format such as HL7's FHIR (Bender and Sartipi, 2013).

In terms of argumentation, there are several directions for future work. Firstly, we will aim to account for various types of interactions that accompany the TMR model, including those concerning conflicts such as side-effects, but also other interactions such as safety. We will also study the integration of preferences from various sources, and possible interactions of those preferences. In addition, we will explore various ways of extracting explanations from the argumentative reasoning process, such as visualising arguments and their relationships, as well as using natural language generation to yield textual explanations of the reasoning outcomes. Finally, we will make use of the well-established theoretical properties of, particularly, ABA^+, regarding reasoning and preferences, and establish what they mean in the context of medical decision making.

Acknowledgements The authors were funded by the EPSRC project **EP/P029558/1** ROAD2H: Resource Optimisation, Argumentation, Decision Support and Knowledge Transfer to Create Value via Learning Health Systems, except for Martin Chapman who was funded by the EPSRC project **EP/P010105/1** CONSULT: Collaborative Mo-

bile Decision Support for Managing Multiple Morbidities.

References

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards Artificial Argumentation. *AI Magazine* 38(3):25–36.

Bao, Z.; Čyras, K.; and Toni, F. 2017. ABAplus: Attack Reversal in Abstract and Structured Argumentation with Preferences. In An, B.; Bazzan, A. L. C.; Leite, J.; Villata, S.; and van der Torre, L., eds., *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference*, Lecture Notes in Computer Science, 420–437. Nice: Springer.

Bender, D., and Sartipi, K. 2013. HL7 FHIR: An Agile and RESTful approach to healthcare information exchange. In *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems*, 326–331. IEEE.

Besnard, P.; García, A. J.; Hunter, A.; Modgil, S.; Prakken, H.; Simari, G. R.; and Toni, F. 2014. Introduction to Structured Argumentation. *Argument & Computation* 5(1):1–4.

Bondarenko, A.; Dung, P. M.; Kowalski, R.; and Toni, F. 1997. An Abstract, Argumentation-Theoretic Approach to Default Reasoning. *Artificial Intelligence* 93(97):63–101.

Čyras, K., and Toni, F. 2016. ABA+: Assumption-Based Argumentation with Preferences. In Baral, C.; Delgrande, J. P.; and Wolter, F., eds., *Principles of Knowledge Representation and Reasoning, 15th International Conference*, 553–556. Cape Town: AAAI Press.

Delaney, B.; Curcin, V.; Andreasson, A.; Arvanitis, T.; Bastiaens, H.; Corrigan, D.; Ethier, J.; Kostopoulou, O.; Kuchinke, W.; McGilchrist, M.; Royen, P.; and Wagner, P. 2015. Translational medicine and patient safety in europe: TRANSFoRm - architecture for the learning health system in europe. *BioMed Research International* 2015:1–8.

Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-person Games. *Artificial Intelligence* 77:321–357.

Fox, J.; Black, L.; Glasspool, D.; Modgil, S.; Oettinger, A.; Patkar, V.; and Williams, M. 2006. Towards a general model for argumentation services. In *Argumentation for Consumers of Healthcare, Papers from the 2006 AAAI Spring Symposium*, 52–57. Stanford: AAAI.

Fraccaro, P.; Arguello Casteleiro, M.; Ainsworth, J.; and Buchan, I. 2015. Adoption of Clinical Decision Support in Multimorbidity: A Systematic Review. *JMIR Medical Informatics* 3(1):e4.

Friedman, C.; Rubin, J.; Brown, J.; Buntin, M.; Corn, M.; Etheredge, L.; Gunter, C.; Musen, M.; Platt, R.; Stead, W.; Sullivan, K.; and Van Houweling, D. 2015. Toward a science of learning systems: a research agenda for the high-functioning learning health system. *Journal of the American Medical Informatics Association* 22(1):43–50.

GOLD. 2017. *Global Strategy for the Diagnosis, Management and Prevention of COPD*. Global Initiative for Chronic Obstructive Lung Disease.

Grace, A.; Mahony, C.; O'donoghue, J.; Heffernan, T.; Molony, D.; and Carroll, T. 2013. A vision for enhancing multimorbid care using clinical decision support systems. *Studies in Health Technology and Informatics* 192(1-2):1117.

Grando, M. A.; Glasspool, D.; and Boxwala, A. A. 2012. Argumentation logic for the flexible enactment of goal-based medical guidelines. *Journal of Biomedical Informatics* 45(5):938–949.

Hunter, A., and Williams, M. 2012. Aggregating evidence about the positive and negative effects of treatments. *Artificial Intelligence in Medicine* 56(3):173–190.

Kokciyan, N.; Sassoon, I.; Young, A.; Chapman, M.; Porat, T.; Ashworth, M.; Curcin, V.; Modgil, S.; Parsons, S.; and Sklar, E. 2018. Towards an Argumentation System for Supporting Patients in Self-Managing their Chronic Conditions. In *Joint Workshop on Health Intelligence (W3PHIAI)*.

Kostopoulou, O.; Porat, T.; Corrigan, D.; Mahmoud, S.; and Delaney, B. C. 2017. Diagnostic accuracy of gps when using an early-intervention decision support system: a high-fidelity simulation. *British Journal of General Practice* 67(656):e201–e208.

Leonardi, G.; Bottrighi, A.; Galliani, G.; Terenziani, P.; Messina, A.; and Della Corte, F. 2012. Exceptions handling within GLARE clinical guideline framework. In *AMIA Annual Symposium Proceedings*, volume 2012, 512–21.

Longo, L. 2016. Argumentation for Knowledge Representation, Conflict Resolution, Defeasible Inference and Its Integration with Machine Learning. In Holzinger, A., ed., *Machine Learning for Health Informatics - State-of-the-Art and Future Challenges*, volume 9605. Springer. 183–208.

McGinnis, J.; Powers, B.; and Grossmann, C. 2011. *Digital Infrastructure for the Learning Health System: The Foundation for Continuous Improvement in Health and Health Care: Workshop Series Summary*. Learning Health System: Workshop Series Summary. National Academies Press.

McGinnis, J. M. 2010. Evidence-based medicine - engineering the learning healthcare system. *Studies in health technology and informatics* 153:145—157.

NICE. 2010. *Chronic obstructive pulmonary disease in over 16s: diagnosis and management*. National Institute for Health and Care Excellence.

Peleg, M. 2013. Computer-interpretable clinical guidelines: A methodological review. *Journal of Biomedical Informatics* 46(4):744–763.

Qassas, M. A.; Fogli, D.; Giacomin, M.; and Guida, G. 2016. ArgMed: A Support System for Medical Decision Making Based on the Analysis of Clinical Discussions. In Papathanasiou, J.; Ploskas, N.; and Linden, I., eds., *Real-World Decision Support Systems: Case Studies*. Springer. 15–41.

Shalom, E.; Shahar, Y.; and Lunenfeld, E. 2016. An architecture for a continuous, user-driven, and data-driven application of clinical guidelines and its evaluation. *Journal of Biomedical Informatics* 59(November):130–148.

Tolchinsky, P.; Cortés, U.; Modgil, S.; Caballero, F.; and López-Navidad, A. 2006. Increasing Human-Organ Transplant Availability: Argumentation-Based Agent Deliberation. *IEEE Intelligent Systems* 21(6):30–37.

Walton, D. 1996. *Argumentation Schemes for Presumptive Reasoning*. L. Erlbaum Associates.

Wilk, S.; Michalowski, M.; Michalowski, W.; Rosu, D.; Carrier, M.; and Kezadri-Hamiaz, M. 2017. Comprehensive mitigation framework for concurrent application of multiple clinical practice guidelines. *Journal of Biomedical Informatics* 66:52–71.

Zamborlini, V.; da Silveira, M.; Pruski, C.; ten Teije, A.; Geleijn, E.; van der Leeden, M.; Stuiver, M.; and van Harmelen, F. 2017. Analyzing interactions on combining multiple clinical guidelines. *Artificial Intelligence in Medicine* 81:78–93.