

Towards an Argumentation System for Supporting Patients in Self-Managing their Chronic Conditions

Nadin Kökciyan,¹ Isabel Sassoon,¹ Anthony P. Young,¹ Martin Chapman,² Talya Porat,²
Mark Ashworth,² Vasa Curcin,^{1,2} Sanjay Modgil,¹ Simon Parsons,¹ Elizabeth Sklar¹

¹ Department of Informatics, King's College London

² Division of Health and Social Care Research, King's College London

Abstract

CONSULT is a decision-support framework designed to help patients self-manage chronic conditions and adhere to agreed-upon treatment plans, in collaboration with healthcare professionals. The approach taken employs *computational argumentation*, a logic-based methodology that provides a formal means for reasoning with evidence by substantiating claims for and against particular conclusions. This paper outlines the architecture of *CONSULT*, illustrating how facts are gathered about the patient and various preferences of the patient and the clinician(s) involved. A logic-based representation of official treatment guidelines by various public health agencies is presented. Logical arguments are constructed from these facts and guidelines; these arguments are analysed to resolve inconsistencies concerning various treatment options and patient/clinician preferences. The claims of the justified arguments are the decisions recommended by *CONSULT*. A clinical example is presented which illustrates the use of *CONSULT* within the context of blood pressure management for secondary stroke prevention.

1 Introduction

Many countries, such as the United Kingdom (UK), have growing populations and comprehensive healthcare systems. Modern improvements in medical diagnosis mean that more people living with multiple chronic morbidities are aware of their conditions. These conditions require constant management by clinicians and thus consume considerable public health resources (Guzman-Castillo et al. 2017). Patients who self-manage their conditions take pressure off public health resources and experience long-term health benefits (Tattersall 2002). As technology has advanced, smartphone and wellness sensor technologies are now capable of recording personal health and activity data that may be relevant for self-management, for example, wellbeing determination in the elderly (Suryadevara and Mukhopadhyay 2012). However, such data may be noisy, and alternative treatment plans can be conflicting. Both patients and clinicians will need to select amongst various treatment plans while also considering issues such as the side effects of drugs, personal treatment preferences and lifestyle constraints.

The *Collaborative Mobile Decision Support for Managing Multiple Morbidities* project¹ seeks to design, verify and implement *CONSULT*, a framework that will gather data from wellness sensors, a patient's own electronic health record (EHR), official clinical guidelines, and input from the patient and their team of carers. *CONSULT* will then use *computational argumentation* to reason with this data, and so justify potential courses of action.

Computational argumentation (Rahwan and Simari 2009), a well-founded logic methodology with roots in philosophy, has been applied in *artificial intelligence* (AI) and *multi-agent systems* as a structured technique for reasoning in which *conclusions* are drawn from *evidence* that *supports* the conclusions. The amenability and the transparency of computational argumentation to human understanding have led to its extensive application in medical decision support systems (Glasspool et al. 2006). As a proof-of-concept, *CONSULT* focuses on the use case of secondary stroke prevention in recovering stroke patients, an important aspect of which is through managing of the patient's blood pressure (BP). *CONSULT* aims to support stroke patients in self-managing their BP, with periodic feedback from clinicians. The reasoning processes articulated in this paper will form a key part towards achieving *CONSULT*'s goals.

Consider Example 1 where a treatment should be offered to a recovering stroke patient. In this paper, we aim to formally represent the knowledge in case studies such as the one in Example 1, and reason with it in order to justify possible treatment plans.

Example 1. *Eric is a 52-year-old male who had just suffered a stroke. He is overweight and has hypertension. When Eric sees his general practitioner (GP), their aim is to prevent Eric from suffering another stroke. It is therefore crucial to keep Eric's BP under control. However, there are several treatment options for the GP to consider, with choice dependent on the priorities of the GP and patient, which may not be aligned. Here, Eric prefers lifestyle changes over drugs but the GP prefers prescribing a drug.*

The contributions of this paper are as follows: (i) We propose a logic-based representation of official treatment guidelines relevant to the treatment of hypertension, (ii) We construct arguments for various treatment options by introduc-

¹UK EPSRC grant EP/P010105/1, 2017–2020.

ing a new argument scheme with associated critical questions, (iii) We use *extended argumentation frameworks* to provide concrete arguments recommending possible courses of treatment given patient data and preferences of both patient and clinician(s).

The outline of this paper is as follows: Section 2 provides background on computational argumentation and argument schemes. In Section 3, we outline the architecture of the CONSULT framework. In Section 4, we illustrate and demonstrate the applicability of our approach on Example 1. In Section 5, we briefly discuss related work. Section 6 highlights directions for future research, such as the treatment of other conditions, multiple-morbidities and polypharmacy, and the data-driven aspects of wellness sensors, in addition to some further challenges for the application of argumentation theory to the medical domain.

2 Background

In this section, we provide an overview of key concepts in computational argumentation (hereafter “argumentation”) theory that are relevant to reasoning about courses of treatment and their possible side effects, given facts about the patient and preferences of the patient and clinicians involved.

Abstract and Extended Argumentation Theory

Argumentation theory is a branch of AI that studies reasoning with incomplete and conflicting information, one application of which is in the field of medical decision support systems. Our starting point is Dung’s abstract argumentation theory (Dung 1995). Arguments are represented with a directed graph $\langle Arg, \mathcal{R} \rangle$ called an (abstract) argumentation framework (AF), where Arg is the set of arguments and $\mathcal{R} \subseteq Arg^2$ is the binary attack relation such that for arguments A and B , $(A, B) \in \mathcal{R}$ iff A attacks (i.e., is a counter-argument to) B . Let $S \subseteq Arg$ be a set of arguments. We say S is *conflict-free* (cf) iff $S^2 \cap \mathcal{R} = \emptyset$; i.e., no two arguments in S attack each other. We say an argument $A \in Arg$ is *acceptable* w.r.t. S iff all attackers of A are in turn attacked by some argument in S . For any S let $d(S) \subseteq Arg$ denote the set of arguments acceptable w.r.t. S . We say S is *self-defending* (sd) iff $S \subseteq d(S)$. We say S is *admissible* iff it is cf and sd. Intuitively, admissible sets of arguments represent justified sets of arguments that are collectively consistent and can respond to all counter-arguments. Since safety is often paramount in medical decision support (Tolchinsky et al. 2012), we use the *grounded extension*, defined to be the \subseteq -smallest admissible set satisfying $S = d(S)$, which always exists and is unique; this captures a conservative form of reasoning where justified arguments are grounded upon incontrovertible truths and are easily computed. Assume that the following dialogue occurs between Eric and his GP:

- GP: “Your test results indicate that you have previously had a mini-stroke.” (Argument A)
- Eric: “Actually, I don’t feel like I had a mini-stroke, therefore I did not have a mini-stroke.” (Argument B)

Abstract argumentation would formalise this as $Arg = \{A, B\}$ and $\mathcal{R} = \{(A, B), (B, A)\}$ because Eric and his

GP disagree. Both $\{A\}$ and $\{B\}$ are admissible sets. However, the grounded extension is \emptyset , so neither $\{A\}$ nor $\{B\}$ is justified. In other words, no recommendation can be made unless we take into account preferences.

Extended argumentation frameworks (EAFs) (Modgil 2009) were developed to enable reasoning about preferences over arguments by incorporating arguments that claim preferences over other arguments. Formally, an EAF is a structure $\langle Arg, \mathcal{R}, \mathcal{D} \rangle$ where $\langle Arg, \mathcal{R} \rangle$ is an AF and $\mathcal{D} \subseteq Arg \times \mathcal{R}$ is the meta-attack relation. If $(X, (A, B)) \in \mathcal{D}$ then this denotes that $X \in Arg$ attacks the attack from A to B by claiming that B is preferred to A , so X would invalidate this attack. Whenever two arguments X and X' express contrary preferences, they would symmetrically attack each other. Formally, if $(X, (A, B))$ and $(X', (B, A)) \in \mathcal{D}$ then (X, X') and $(X', X) \in \mathcal{R}$. The notion of admissible sets can be appropriately generalised to EAFs, but for our purposes, arguments that are not attacked are justified, while arguments that are attacked by justified arguments cannot be justified. Attacks that are meta-attacked by justified arguments are rendered ineffective.

If we assume that the GP’s argument (Argument A) is stronger or preferred to Eric’s argument (Argument B), then we can represent this as a preference argument (Argument C), claiming that A is preferred to B . We represent this by a meta-attack relation $\mathcal{D} = \{(C, (B, A))\}$. $\{B\}$ is no longer justified since C is attacking the attack (B, A) , hence A is justified and the grounded extension is $\{A, C\}$.

Argument Schemes and Critical Questions

In practical reasoning, arguments can be challenged and defeated by further arguments. It is therefore possible to identify more arguments and consider alternatives, if any. Intuitively, arguments should first be challenged, then become justified and taken into consideration if they survive being defeated. One way of doing this is using argument schemes and critical questions.

Argument schemes (Walton, Reed, and Macagno 2008) are semi-formal representations of the structures of common types of arguments. One of the key features of argument schemes is the list of associated *critical questions* (CQs). The claim that a scheme supports is presumptive and the claim is withdrawn unless the CQs posed have been answered successfully. The instantiation of the appropriate argument scheme, in conjunction with its associated CQs is a method of generating a set of arguments. The inference mechanism characterized by the argument scheme will ensure that only arguments that have not been defeated by the CQs will be generated.

Table 1 shows Walton’s *Sufficient Condition Scheme* for practical reasoning. This scheme states that an agent should perform an action if this action helps that agent to achieve its goal. Walton proposes four CQs: (1) Are there alternative ways of realising goal G ?, (2) Is it possible to do action A ?, (3) Does agent a have goals other than G which should be taken into account?, and (4) Are there other consequences of doing action A which should be taken into account? These questions can serve as counterarguments for arguments that conform to the Sufficient Condition Scheme. For example,

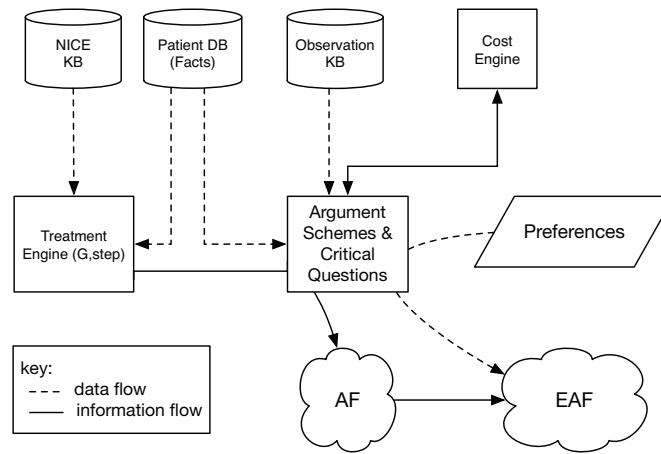


Figure 1: The architecture of the knowledge bases, databases and argument schemes used by CONSULT

according to the first CQ, if there are alternative ways of carrying out the same goal, then these alternatives may change the outcome of the decision process of the agent.

AS
G is a goal for agent a
Doing action A is sufficient for a to carry out goal G
Therefore agent a ought to do action A .

Table 1: Walton’s Sufficient Condition Scheme

3 CONSULT

In this section, we outline the CONSULT framework and explain its various components. Figure 1 illustrates the architecture of CONSULT, enumerating the *knowledge bases* (KB) and *databases* (DB) considered and showing how aspects of argumentation come into play. We distinguish a knowledge base, a set of rules, from a database, a set of facts or data points.

When prescribing a treatment plan for a patient, GPs in the UK follow the *National Institute for Health and Care Excellence* (NICE) guidelines for treatment options, while also taking into account patient-specific data, treatment costs and patient / clinician preferences.

The CONSULT framework aims to support this process by identifying arguments that justify various treatment options specific to the patient. It includes a Patient Database (DB), which contains all available data and facts about each patient. Its NICE Knowledge Base (KB) represents the most relevant clinical guidelines for treating the patient. CONSULT constructs arguments for and against various treatments specific to a patient with its Treatment Engine, which uses argument schemes that are subjected to CQs. One of the CQs leverages the Cost Engine to ensure that any treatment cost considerations are applied, hence arguments for equally effective but more expensive treatments will be defeated. The arguments not defeated by the CQs will form an AF. One aim is for CONSULT to be able to reason with and about GP and patient preferences through the use of EAFs.

Step (1.a):

“Offer people aged under 55 years step 1 antihypertensive treatment with an ACE inhibitor or a low-cost ARB.

If an ACE inhibitor is prescribed and is not tolerated (for example, because of cough), offer a low-cost ARB. Beta-blockers are not a preferred initial therapy for hypertension. However, beta-blockers may be considered in younger people, particularly: (1) those with an intolerance or contraindication to ACE inhibitors and angiotensin II receptor antagonists, or (2) women of child-bearing potential, or (3) people with evidence of increased sympathetic drive.”

Step (1.b):

“Offer step 1 antihypertensive treatment with a CCB to people aged over 55 years and to black people of African or Caribbean family origin of any age.

If a CCB is not suitable, for example because of oedema or intolerance, or if there is evidence of heart failure or a high risk of heart failure, offer a thiazide-like diuretic.”

Table 2: Step 1 of NICE guideline CG127

Knowledge Representation for the Treatment of Hypertension

In order for CONSULT to reason about treatment plans, we represent knowledge in the hypertension domain using first order logic. For example, we represent the hypertension treatment guideline CG127 published by NICE (NICE 2016) (see Table 2). Patient characteristics, such as ethnicity or experienced side effects could change the treatment plan. We then represent the treatment options for hypertension by following the patient information leaflet provided by the UK’s National Health Service (NHS) Choices.

Representation of the Relevant NICE Guideline NICE has a set of guidelines to help healthcare professionals in diagnosing and treating primary hypertension, and thereby reducing the risk of primary and secondary strokes. The guideline CG127 mentions four types of drugs: A, B, C and D. A refers to *ACE Inhibitor* or *low-cost Angiotensin II receptor blocker* (ARB), B refers to *Beta-blocker*, C refers to *calcium-channel blocker* (CCB) and D refers to *thiazide-like Diuretic*. The guideline includes treatment steps, such that a patient progresses to the next step and takes a new drug if their BP does not improve in the previous step. The guideline provides guidance on which of the treatments or treatment combinations should be considered at each step.

Step (1.a):

$(age < 55) \rightarrow offer(A_1, S_1, d) \vee offer(A_2, S_1, d)$
 $\neg tolerated(A_1) \rightarrow offer(A_2, S_1, d) \wedge \neg offer(A_1, S_1, d)$
 $\neg tolerated(A_2) \rightarrow \neg offer(A_2, S_1, d)$
 $\neg tolerated(A_1) \vee \neg tolerated(A_2) \rightarrow offer(B, S_1, d)$
 $chbearing\text{-}potential \vee inc\text{-}sympa\text{-}drive \rightarrow offer(B, S_1, d)$

Step (1.b):

$(age \geq 55) \vee bl\text{-}afr \vee bl\text{-}car \rightarrow offer(C, S_1, d)$
 $\neg tolerated(C) \rightarrow \neg offer(C, S_1, d) \wedge offer(D, S_1, d)$
 $oedema \vee heart\text{-}failure \vee hr\text{-}heart\text{-}failure \rightarrow offer(D, S_1, d)$

All Steps:

$\rightarrow offer(LS, Y, -)$
 $offer(A_1, Y, d) \rightarrow \neg offer(A_2, Y, d)$
 $offer(A_2, Y, d) \rightarrow \neg offer(A_1, Y, d)$
 $offer(X, Y, high\text{-}dose) \rightarrow \neg offer(X, Y, low\text{-}dose)$
 $offer(X, Y, low\text{-}dose) \rightarrow \neg offer(X, Y, high\text{-}dose)$

Table 3: Step 1 anti-hypertensive treatment as logic rules

To represent CG127 formally, we denote each treatment step as S_i , which represents the i -th step in the treatment plan. Table 2 is the guideline for S_1 . Let A_1 denote ACE Inhibitor and A_2 denote ARB. We formally represent this information using the logic rules shown in Table 3. The other treatment steps are represented formally in the same manner. The guideline and the rules of these steps are shown in Tables 4 and 5, respectively. Note that, for simplicity, we only represent part of the NICE guideline. For example, we do not indicate either thiazide-like diuretic names (e.g., chlortalidone) or drug dosages (e.g., 12.5–25.0 mg once daily).

Each rule is of the form $P \rightarrow Q$, which means that if the antecedent P holds, then the consequent Q also holds. Both P and Q can consist of disjunctions (\vee) and conjunctions (\wedge) of atoms. The atoms are the facts about the patient (e.g., Black-African). During the treatment, the GP can decide to use different doses of a drug such as the maximally tolerated dose (high-dose) or the minimal effective dose (low-dose). The atom $offer(X, Y, d)$ states that the drug $X \in \{A_1, A_2, B, C, D, \alpha B\}$ (where αB denotes alpha-blocker) should be prescribed in Step $Y \in \{S_i\}_{i=1}^4$ with dose d (high-dose or low-dose). For example, in Step (1.a), a white male patient aged 40 who is intolerant to A_1 can be offered low-dose A_2 or B .

The NICE guideline points out that under some condi-

Step 2:

“If blood pressure is not controlled by step 1 treatment, offer step 2 treatment with a CCB in combination with either an ACE inhibitor or an ARB.

If a CCB is not suitable for step 2 treatment, for example because of oedema or intolerance, or if there is evidence of heart failure or a high risk of heart failure, offer a thiazide-like diuretic.

For black people of African or Caribbean family origin, consider an ARB in preference to an ACE inhibitor, in combination with a CCB.

If therapy is initiated with a beta-blocker and a second drug is required, add a CCB rather than a thiazide-like diuretic to reduce the person’s risk of developing diabetes.”

Step 3:

“If treatment with three drugs is required, the combination of ACE inhibitor or ARB, CCB and thiazide-like diuretic should be used.”

Step 4:

“Consider further diuretic therapy.

If further diuretic therapy for resistant hypertension at step 4 is not tolerated, or is contraindicated or ineffective, consider an alpha- or beta-blocker.”

Table 4: Steps 2-4 of NICE guideline CG127

tions, it is better to choose one treatment over another. For example, if a patient uses a beta-blocker (B) in his therapy and a second drug is required, then it is better to offer CCB (C) to reduce the patient’s risk of developing diabetes (Step 2 in Table 4). We represent such preferences by the atom $pref(Y, Z)$, where Y and Z are possible treatment options, which states that Z is a more preferred treatment than Y . For example:

$$pref(offer(D, S_2, d), offer(C, S_2, d))$$

represents such a preference from the NICE guideline.

For all of the treatment steps in the NICE guideline, A_1 and A_2 cannot be used together, and a treatment can only be used in a single dose (either a low-dose or a high-dose). These restrictions are also defined as logic rules, as shown at the bottom of Table 3. Moreover, the GP has the option to treat hypertension with lifestyle changes (e.g., losing weight, eating a healthy diet and exercise); we represent such treatment options as LS . In future work, CONSULT will be able to monitor a patient’s actual lifestyle changes; e.g., through wellness sensor measurements of daily activity and weight.

Representation of the NHS Choices Leaflet Each treatment requires use of drugs that may result in negative side effects. In such cases, healthcare professionals may try alternative treatments. In Table 6, we show how observing various side effects affects treatment options (NHS 2016), represented as rules in first order logic. For example, if there

Step 2:

$$\begin{aligned} & \text{offer}(A_1, S_1, d) \vee \text{offer}(A_2, S_1, d) \rightarrow \text{offer}(C, S_2, d) \\ & \text{offer}(C, S_1, d) \rightarrow \text{offer}(A_1, S_2, d) \vee \text{offer}(A_2, S_2, d) \\ & \neg \text{tolerated}(C) \rightarrow \neg \text{offer}(C, S_2, d) \wedge \text{offer}(D, S_2, d) \\ & \text{oedema} \vee \text{heart-failure} \vee \text{hr-heart-failure} \rightarrow \text{offer}(D, S_2, d) \\ & \text{bl-afr} \vee \text{bl-car} \rightarrow \text{offer}(A_1, S_2, d) \vee \text{offer}(A_2, S_2, d) \\ & \text{bl-afr} \vee \text{bl-car} \rightarrow \text{pref}(\text{offer}(A_1, S_2, d), \text{offer}(A_2, S_2, d)) \\ & \text{offer}(B, S_1, d) \rightarrow \text{offer}(C, S_2, d) \vee \text{offer}(D, S_2, d) \\ & \text{offer}(B, S_1, d) \rightarrow \text{pref}(\text{offer}(D, S_2, d), \text{offer}(C, S_2, d)) \end{aligned}$$
Step 3:

$$\rightarrow \text{offer}(D, S_3, d)$$
Step 4:

$$\begin{aligned} & \rightarrow \text{offer}(D, S_4, d) \\ & \neg \text{tolerated}(D) \rightarrow \text{offer}(\alpha B, S_4, d) \vee \text{offer}(B, S_4, d) \end{aligned}$$

Table 5: Steps 2-4 anti-hypertensive treatment as logic rules

is evidence of flu-like symptoms, then the GP can prescribe A_2 instead of A_1 during the treatment process.

pregnancy \vee breastfeeding \rightarrow $\neg \text{offer}(A_1, S, d) \wedge \neg \text{offer}(A_2, S, d) \wedge \neg \text{offer}(D, S, d)$
dry-cough \vee dizziness \vee headaches \vee rash \rightarrow $\neg \text{offer}(A_1, S, d)$
dizziness \vee headaches \vee flu-like-symptoms \rightarrow $\neg \text{offer}(A_2, S, d)$
headaches \vee swollen-ankles \vee constipation \rightarrow $\neg \text{offer}(C, S, d)$
dizziness \vee increased-thirst \vee increased-toilet-frequency \vee rash $\rightarrow \neg \text{offer}(D, S, d)$
erectile-dysfunction \vee fall-in-potassium-levels \rightarrow $\neg \text{offer}(D, S, d)$
dizziness \vee headaches \vee tiredness \vee cold-hands-feet \rightarrow $\neg \text{offer}(B, S, d)$
dizziness \vee light-headedness \vee fainting \rightarrow $\neg \text{offer}(\alpha B, S, d)$

Table 6: Anti-hypertensive Treatment Options as defined by NHS. S is the current treatment step for the patient.**Argument Scheme for Proposed Treatment**

In order to generate arguments in support of different treatment options, we use an argument scheme structure similar to the practical reasoning scheme (Walton, Reed, and Macagno 2008). Our argument scheme generates an argument in support of each possible treatment, given the known Facts F about the patient and the treatment goal G to be realised, e.g., lowering the patient's BP. The arguments instantiated by this scheme are all subject to CQs. In this case the critical questions are used to generate counterarguments to the arguments instantiated by the AS. These counterarguments will be generated when a treatment has either been used before unsuccessfully or has caused side effects, as well when an equivalent cheaper treatment is possible.

In Table 7, we propose an argument scheme – the argument scheme for a Proposed Treatment (ASPT). The patient facts F include their age, BP (including stage), ethnicity, previous treatments and the current treatment step. The Treatment Engine reasons with the patient facts to find possible treatments from the NICE KB. Note that each possible treatment conforms to the ASPT.

ASPT
<i>premise</i> - Given the patient Facts F
<i>premise</i> - In order to realise the goal G
<i>premise</i> - Treatment T promotes the goal G
therefore : Treatment T should be considered

Table 7: Argument scheme for a proposed treatment

Critical Questions. Arguments instantiating the scheme will be subject to the following CQs:

CQ 1. Has this treatment been unsuccessfully used on the patient in the past?

CQ 2. Has the patient experienced side effects from this treatment in the past?

CQ 3. Is there an equivalent cheaper treatment for the treatment step of the patient?

If ASPT yields an argument in support of treatment T , then the CQs have the potential to attack or yield additional arguments for possible treatments. For example, if the Cost Engine indicates an equivalent cheaper treatment, then CQ3 yields a counter-argument proposing this treatment.

The resulting set of arguments will form an AF. An example is illustrated in Figure 2, where the argument framework (AF) consists of three arguments and four attacks between arguments. We will explore this AF in more detail in the next section.




Figure 2: The Argument Framework for Example 1

Reasoning with Preferences

There are differing orders of preferences over the possible treatments for a patient, namely ones from the GP (who may prefer the most effective treatment) and ones from the patient (who may prefer to minimise side effects). There may also be additional preference orders from external sources such as secondary care specialists, other GPs or the patient's family members who are involved in managing their care.

In order to derive the meta-level arguments, these preference orders over treatments need to be expressed as attacks

over the attacks between arguments in support of treatments. For example, if we have the expressed preference between the treatments x and y such that $x \succ y$ when x is strictly more preferred, then this can be represented as an argument $(i, (y, x))$ where i denotes the argument for preferring x to y . Intuitively, if x is preferred to y , then the preference argument i attacks the attack from y to x , as y is less preferred.

The meta-level arguments expressing the preference orders will form part of the EAF. Different sets of preferences can be considered simultaneously in an EAF by deciding a priority between the different preference orders. A treatment argument is justified if it is part of the grounded extension of the EAF.

4 Execution of the Running Example

We now return to our sample patient scenario to illustrate our system, introduced in Example 1. The facts about Eric from the Patient DB are formally represented as follows: $f_{eric} = \{\text{age}=52, \text{ethnicity}=\text{white}, \text{overweight}\}$. Eric has never been prescribed medication for hypertension before, as such he is in step one (S_1) and should be offered only one treatment. By instantiating the argument scheme ASPT, as shown in Table 8, the Treatment Engine generates arguments each in support of one of five treatments that are shown in Table 9.

ASPT(Eric, t_i)
<i>premise</i> - Given the patient facts f_{eric}
<i>premise</i> - In order to realise the goal G
<i>premise</i> - Treatment t_i promotes the goal G
therefore : Treatment t_i should be considered

Table 8: Instantiating ASPT for Eric per treatment

However, these arguments for possible treatments are subject to CQs, as follows. As Eric has not been prescribed any BP medication in the past, he did not experience any side effects from such a treatment (CQ 2) and no BP treatments are known to have been unsuccessful for Eric (CQ 1). Only CQ 3 instantiates counter-arguments in Example 1: t_2^l and t_2^h are more expensive but their treatment outcomes are equivalent to t_1^l and t_1^h , respectively. This information is provided by the Cost Engine. Accordingly, t_2^l and t_2^h are defeated since there are cheaper treatment options. The set of arguments in support of possible treatments is reduced to three arguments (t_{ls} , t_1^l and t_1^h), each in support of a different treatment. These arguments are added to the AF since they conform to the CQs (CQ 1, CQ 2 and CQ 3).

Figure 2 depicts the current AF. Each argument in support of a treatment is represented as a node. The arrows denote the attacks between the arguments. Argument t_1^l and t_1^h are attacking each other as these treatments cannot be offered together (Table 3). Eric views t_{ls} as an alternative treatment, which is mutually exclusive to both t_1^h and t_1^l ; hence there are asymmetric attacks between t_{ls} and the other treatment arguments in the AF.

A treatment should be chosen by considering the preferences over treatments. In Example 1, there are two sets of preferences: the patient’s (Eric) and the GP’s. Eric prefers

t_{ls} : offer(LS , S_1 , -)
t_1^l : offer(A_1 , S_1 , low-dose)
t_1^h : offer(A_1 , S_1 , high-dose)
t_2^l : offer(A_2 , S_1 , low-dose)
t_2^h : offer(A_2 , S_1 , high-dose)

Table 9: Possible Treatments for Eric as recommended by the Treatment Engine

making lifestyle changes; i.e., $t_{ls} \succ t_1^h$ and $t_{ls} \succ t_1^l$. The GP prefers prescribing some drug; i.e., $t_{ls} \prec t_1^h$ and $t_{ls} \prec t_1^l$, and the GP may prefer to treat with the higher tolerated dose; i.e., $t_1^h \succ t_1^l$. The meta-level arguments are derived from the preference relations as follows:

- For Eric: $\{(e_l, (t_1^l, t_{ls})), (e_h, (t_1^h, t_{ls}))\}$
- For the GP: $\{(g_h, (t_{ls}, t_1^h)), (g_l, (t_{ls}, t_1^l)), (g_d, (t_1^h, t_1^l))\}$

The preference orders resulting from Eric’s and the GP’s preferences are also in direct conflict, therefore there are additional attacks between these. These are illustrated in the EAFs in Figure 3, where each EAF displays different precedences between the preferences of Eric and the preferences of the GP.

Figure 3a illustrates the EAF resulting from Eric’s preferences taking precedence. The grounded extension in this EAF is $\{t_{ls}, g_d, e_h, e_l, (e_h \succ g_h), (e_l \succ g_l)\}$. Hence, the only treatment argument in the grounded extension is t_{ls} . In Figure 3b, the GP’s preferences take precedence over Eric’s. In the resulting EAF, the grounded extension is $\{t_1^h, t_{ls}, g_d, g_l, (g_h \succ e_h), (g_l \succ e_l), g_h\}$. This set contains both t_1^h and t_{ls} , so these two treatments are justified in this setting.

Should the GP want to explicitly exclude lifestyle changes from the set of possible treatments, then this would be achieved by an argument $\neg t_{ls}$ that would attack t_{ls} . This could be a relevant option if the patient’s physical condition would not allow sufficient changes to affect BP.

5 Related Work

Over the last few decades, argumentation theory has been applied to a range of subjects including multi-agent systems, game theory, legal reasoning and machine learning (Rahwan and Simari 2009). In the medical domain, argumentation theory has been applied to medical expert systems to make recommendations with clear reasons supporting them based on the data given, and whether a given course of treatment is safe to administer (Fox, Glasspool, and Bury 2001; Fox et al. 2007; Glasspool et al. 2007). CONSULT aims to be able to explain the recommendations it gives in a similar way, but has the additional ability to reason about patients’ and clinicians’ preferences using EAFs. Further, Hunter and Williams have proposed argumentation-based techniques to aggregate the conclusions of various clinical trials to determine which of two treatments is more effective given the situation (Hunter and Williams 2010). These techniques can be useful to CONSULT if we consider incorporating the latest

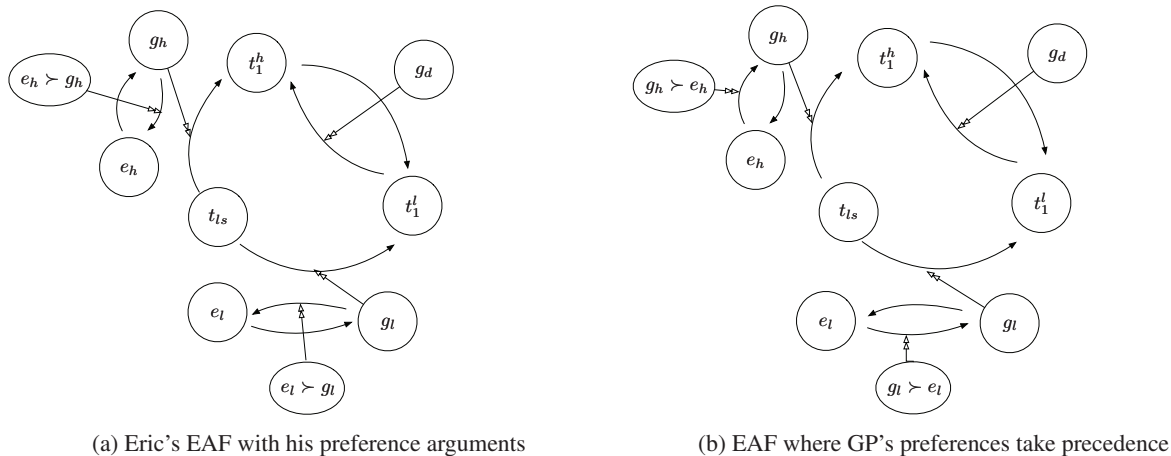


Figure 3: EAFs according to the perspectives of Eric and his GP

clinical trials relevant to its various treatment recommendations.

Atkinson *et al.* propose an argumentation-based approach for reasoning with defeasible arguments (Atkinson, Bench-Capon, and Modgil 2006). To show the applicability of their approach, they model a DRAMA (Deliberative Reasoning with Arguments about Actions) agent that would recommend a treatment based on arguments collected from various information sources. Similar to CONSULT, they make use of argument schemes and multiple knowledge bases. In their model, each argument is associated with values such as safety and efficacy. Hence, they recommend treatments with higher values regarding a strict partial ordering on the values. Our work differs from theirs in that we consider a strict partial ordering on the arguments and use EAFs for defeasible reasoning where meta-level attacks are also possible.

Reasoning with arguments that are collected from various information sources is a challenging problem since each information source is of varying trustworthiness. The *ArgTrust* framework was developed (Tang, Sklar, and Parsons 2012) and evaluated (Sklar et al. 2016) as a decision-support tool in which the evidence that influences a recommendation is modulated according to values of trust that a user places on the evidence. In their work, they introduce a formal argumentation system for reasoning with the collected information. Similar to their work, we would like to associate arguments with trust values depending on the arguments' sources of data in order to recommend more reliable treatments.

6 Summary and Future Work

In this paper, we have introduced the argumentation-based decision support system *CONSULT*, which aims to assist healthcare professionals in choosing treatments for their patients, as well as patients in self-managing their chronic conditions. We have illustrated how *CONSULT* is designed to work in the context of treating high blood pressure in recovering stroke patients. We have provided a formal representation of the NICE guideline CG127 and the NHS Leaflet for hypertension treatment options. In our proposed approach,

the Treatment Engine instantiates possible treatment arguments given patient information using the ASPT argument scheme and subjects these to critical questions. As a result of this step, some arguments are defeated if they do not conform with ASPT's critical questions. Meta-level arguments are generated from the preferences expressed by the patient and the GP. An extended argumentation framework (EAF) is generated from the treatment arguments and the meta-level preference arguments. The grounded extension of the EAF is computed by considering different precedences between the sets of preferences. The presence of a treatment argument in a grounded extension justifies it as a treatment to recommend in the given circumstances. We have illustrated the applicability of our approach through a running example.

In ongoing work, we are implementing and evaluating the components of the *CONSULT* architecture outlined here. The aim is to deploy *CONSULT* on a mobile device such as a tablet, with intuitive dashboards for clinicians and patients. We are planning to evaluate on more complex scenarios, through focus groups and user studies. *CONSULT* design and features are currently being informed by focus group interviews consisting of recovering stroke patients and their carers. In addition to patient facts and patient / clinician preferences, we will automatically construct arguments from patient data obtained through commercial wellness sensors, the patient's electronic health record, and extensive clinical guidelines automatically extracted off the NHS or NICE websites; this will inform personalised treatment plans. Techniques that track data provenance will be employed to help determine the priority and trustworthiness of such data, and hence the importance of each argument constructed, which in turn will help determine which arguments are justified (Modgil and Prakken 2010; 2013). Here, we have considered how *CONSULT* might help a patient self-manage one condition, namely high blood pressure; but in future, argumentation theory will be applied to resolve conflicting treatments in the case of multiple morbidities and related issues in polypharmacy. Further, patients and clinicians will be able to understand why such treat-

ments are recommended by CONSULT through the theory of *dialogical argumentation* (McBurney and Parsons 2009; Modgil 2017), where the reasons for a claim can be explicitly traced back to its supporting facts and how its counter-arguments have been defeated.

As CONSULT extracts and aggregates data about each patient from multiple sources, it will be possible to leverage this data to benchmark a patient's additional risks in order to further personalise treatment. Future work will consider benchmarking models that may be derived from statistical models. In order for CONSULT to be able to exploit these and aggregate all the patient-related conclusions from such models, we will be exploring if and how arguments can be derived from the statistical models of the data. Furthermore, we will explore how these quantitative arguments from the models can be considered alongside the qualitative arguments such as the ones generated from clinical guidelines, such as the ones in this paper. This will necessitate further work in the theory of relating argument schemes and critical questions to reasoning about preferences between other arguments, and the implementation of argumentation engines that can reason with preferences and numerical data.

Acknowledgements

The CONSULT project is supported by the UK Engineering & Physical Sciences Research Council (EPSRC) under grant #EP/P010105/1.

References

- Atkinson, K.; Bench-Capon, T. J. M.; and Modgil, S. 2006. Argumentation for decision support. In *Database and Expert Systems Applications (DEXA)*, 822–831.
- 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(2):321–358.
- Fox, J.; Glasspool, D.; Grecu, D.; Modgil, S.; South, M.; and Patkar, V. 2007. Argumentation-based inference and decision making—A medical perspective. *IEEE Intelligent Systems* 22(6).
- Fox, J.; Glasspool, D.; and Bury, J. 2001. *Quantitative and Qualitative Approaches to Reasoning Under Uncertainty in Medical Decision Making*. Springer. 272–282.
- Glasspool, D.; Fox, J.; Oettinger, A.; and Smith-Spark, J. 2006. Argumentation in Decision Support for Medical Care Planning for Patients and Clinicians. In *AAAI Spring Symposium: Argumentation for Consumers of Healthcare*, 58–63.
- Glasspool, D.; Oettinger, A.; Smith-Spark, J.; Castillo, F.; Monaghan, V.; Fox, J.; et al. 2007. Supporting medical planning by mitigating cognitive load. *Methods of Information in Medicine* 46(6):636–640.
- Guzman-Castillo, M.; Ahmadi-Abhari, S.; Bandosz, P.; Capewell, S.; Steptoe, A.; Singh-Manoux, A.; Kivimaki, M.; Shipley, M. J.; Brunner, E. J.; and O'Flaherty, M. 2017. Forecasted trends in disability and life expectancy in England and Wales up to 2025: a modelling study. *The Lancet Public Health*.
- Hunter, A., and Williams, M. 2010. Argumentation for Aggregating Clinical Evidence. In *IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, volume 1, 361–368.
- McBurney, P., and Parsons, S. 2009. *Dialogue Games for Agent Argumentation*. Springer. 261–280.
- Modgil, S., and Prakken, H. 2010. Reasoning about Preferences in Structured Extended Argumentation Frameworks. In *Conference on Computational Models of Argument (COMMA)*, 347–358.
- Modgil, S., and Prakken, H. 2013. A General Account of Argumentation with Preferences. *Artificial Intelligence* 195:361–397.
- Modgil, S. 2009. Reasoning about preferences in argumentation frameworks. *Artificial Intelligence* 173(9–10):901 – 934.
- Modgil, S. 2017. Towards a General Framework for Dialogues that Accommodate Reasoning About Preferences. In *International Workshop on Theory and Applications of Formal Argumentation (TFAFA)*.
- NHS. 2016. Treatment options for high blood pressure (hypertension). [http://www.nhs.uk/Conditions/Blood-pressure-\(high\)/Pages/treatmentoptions.aspx](http://www.nhs.uk/Conditions/Blood-pressure-(high)/Pages/treatmentoptions.aspx). last accessed 17/10/2017.
- NICE. 2016. Hypertension in adults: diagnosis and management. <https://www.nice.org.uk/guidance/cg127>. last accessed 17/10/2017.
- Rahwan, I., and Simari, G. R. 2009. *Argumentation in Artificial Intelligence*. Springer.
- Sklar, E. I.; Parsons, S.; Li, Z.; Salvit, J.; Perumal, S.; Wall, H.; and Mangels, J. 2016. Evaluation of a trust-modulated argumentation-based interactive decision-making tool. *Autonomous Agents and Multi-Agent Systems* 30(1):136–173.
- Suryadevara, N. K., and Mukhopadhyay, S. C. 2012. Wireless Sensor Network Based Home Monitoring System for Wellness Determination of Elderly. *IEEE Sensors Journal* 12(6):1965–1972.
- Tang, Y.; Sklar, E.; and Parsons, S. 2012. An Argumentation Engine: ArgTrust. In *International Workshop on Argumentation in Multiagent Systems (ArgMAS)*.
- Tattersall, R. 2002. The Expert Patient: A New Approach to Chronic Disease Management for the Twenty-First Century. *Clinical Medicine* 2(3):227–229.
- Tolchinsky, P.; Modgil, S.; Atkinson, K.; McBurney, P.; and Cortés, U. 2012. Deliberation Dialogues for Reasoning about Safety Critical Actions. *Autonomous Agents and Multi-Agent Systems* 25(2):209–259.
- Walton, D.; Reed, C.; and Macagno, F. 2008. *Argumentation Schemes*. Cambridge University Press.