

Dialogue Games for Explaining Medication Choices^{*}

Qurat-ul-ain Shaheen^[0000-0003-2258-4476], Alice Toniolo^[0000-0002-6816-6360],
and Juliana K. F. Bowles^[0000-0002-5918-9114]

School of Computer Science, University of St Andrews, St Andrews KY16 9SX, UK
quratulainshaheen@gmail.com
{a.toniolo,jkfb}@st-andrews.ac.uk

Abstract. SMT solvers can be used efficiently to search for optimal paths across multiple graphs when optimising for certain resources. In the medical context, these graphs can represent treatment plans for chronic conditions where the optimal paths across all plans under consideration are the ones which minimize adverse drug interactions. The SMT solvers, however, work as a black-box model and there is a need to justify the optimal plans in a human-friendly way. We aim to fulfill this need by proposing explanatory dialogue protocols based on computational argumentation to increase the understanding and trust of humans interacting with the system. The protocols provide supporting reasons for nodes in a path and also allow counter reasons for the nodes not in the graph, highlighting any potential adverse interactions during the dialogue.

Keywords: Explanations · Dialogues Games · Argumentation · SMT solvers

1 Introduction

Intelligent systems are becoming increasingly popular in today’s digital world with the pervasiveness of Artificial Intelligence (AI). The focus on human-friendly integration of AI has increased in recent years with growing awareness of the need for transparency and justifiability of the black-box recommendations made by some of these systems. Limited transparency and justifiability has been shown to hamper the mainstream adoption of these systems [2].

An even more urgent need for clarity comes with critical systems, such as clinical decision support systems, and there are ongoing efforts towards eXplainable Artificial Intelligence (XAI) [2], specially for medical XAI [23]. While most of the existing approaches to XAI focus on *interpretability* of machine learning models [1], rule-based approaches can be used to *justify* decisions of non-interpretable systems by explaining why a decision is a good one rather than how it was made. We follow the later approach here.

^{*} This research was conducted whilst the first author was visiting the University of St Andrews.

Medical treatment plans (aka clinical guidelines) are drawn from evidence-based recommendations, and published in the UK by the National Institute for Health and Care Excellence (NICE)¹ to document treatment practices for specific chronic diseases such as diabetes, hypertension, and so on. The plans capture a series of steps in the management of the disease, giving an indication of which advice to follow including which drugs (e.g., from a group) to prescribe at a given stage of the disease and what changes to do once the disease progresses (e.g., add a specific medication). These plans, however, fall short for patients with multiple chronic health conditions, known as *multimorbidity*, and in particular they do not account for the potential adverse drug reactions (ADRs) caused by drugs taken for different conditions [14]. Treatment plans for individual chronic conditions can essentially be regarded as graphs, and in the presence of multimorbidities we are in effect trying to search for the optimal path across several graphs which minimises ADRs as much as possible. Satisfiability Modulo Theories (SMT) solvers are popular in analysis and model checking settings because of their expressive power, arithmetic capabilities and scalability [19]. SMT solvers have been used efficiently to identify and minimise ADRs in [5–7]. There is, however, still a need to add clarity to the identified optimal path and justify why one combination of paths across different treatment plans may be preferable to another.

The use of computational argumentation has been popular in healthcare [4, 15] and more recently for managing multimorbidity [9, 20], because its formulation and evaluation of pros and cons can be seen as emulating human critical thinking. Recently argumentation has also generated interest for its explainability potential [12, 13] because of its reasoning transparency. Artificial dialogue based on argumentation allows agents to engage in a gradual exploration of reasons for a conclusion and potential disagreement. In human argumentation, argumentation-based dialogue has shown to be effective in improving understanding of scientific topics [10]. Consequently we think argumentation is a great fit for tackling the problem of justifying optimal paths for multimorbidity treatment plans in a way that ensures human engagement.

The contribution of this paper is a novel approach for justifying recommendations of an SMT solver in an interactive way through argumentation-based dialogues. We use argumentation to augment these recommendations with explanations and propose a novel model of explanatory dialogue protocols to allow human-like engagement with the explanation model. The *explanations* model the underlying reasoning of the SMT solver while the *dialogue protocols* highlight any unresolved issues in the recommended solution.

The remaining paper is organised as follows. Section 2 summarises the background of the argumentation model underlying the explanation dialogues. Section 3 presents the dialogue protocols. Finally, Section 4 concludes and identifies directions for future research.

¹ For details see <https://www.nice.org.uk>

2 Argumentation for pharmaceutical treatment plans

Here we summarise the background research for the work presented in this paper and introduce relevant underlying concepts related to argumentation.

2.1 Finding optimal plans with SMT solver

We base our approach on the work of Kovalov and Bowles [16] which identifies optimal treatment plans for multimorbid patients using the optimising SMT solver Z3 [18]. The authors manually convert the flowchart representations of clinical guidelines for treating specific health conditions into a compact graph which they refer to as a *Pharmaceutical Graph* (PG). A PG is a directed acyclic graph where the root node represents the diagnosed disease, and all the other nodes represent drugs or groups of drugs to be given to a patient. A maximal path in the graph represents a complete treatment plan.

Kovalov and Bowles [16] first create PGs corresponding to different health conditions for a hypothetical multimorbid patient and feed them to the SMT solver to identify the optimal paths for each of the PGs such that adverse drug reactions across all PGs are minimised. They take three types of ADRs into account: drug-drug interactions, drug-disease interactions and drug-patient interactions. Z3 is a Boolean satisfiability problem modulo theories (SMT) solver which finds the optimal assignment given by some objective function [3]. The objective function used by the authors, called *score*, is computed as a combination of medicine efficacy (positive score) and drug interaction conflict (negative score). More comprehensive extensions that incorporate side-effects, time and patient preferences have been recently developed in [7]. The approach is scalable to any number of drug alternatives.

2.2 Modelling Explanations

The resulting maximal path described above is a sequence of drugs or group of drugs for a specific health condition [16]. While PGs encode a score wrt known ADRs to extract a most effective solution, when presenting the sequence of treatments to a health professional or to a patient, the reasons for this to be indeed the most effective solution are not immediately accessible. This means that the user might have to search for additional information to better understand the solution provided. To facilitate this process, we design a multi-layer explanation model in order to have flexible interactions that can adapt to the user's need for clarification. We organise the graph into abstraction layers and explicitly mark the abstractions levels in the graph structure in terms of branches, sub-branches and nodes. We refer to the resulting graph as *PGraph*. We define four abstraction levels for the PGraph and the explanations: *Branch*, *Entity*, *Drug* and *Group*. A *Branch* is a cluster of nodes joined by an initial node and a terminating node such that initial node has only one incoming edge and the terminating node has only one outgoing edge. An *Entity* can be an atomic node representing a *Drug* or a composite node representing a group of drugs (*Group*). A *Group* is a

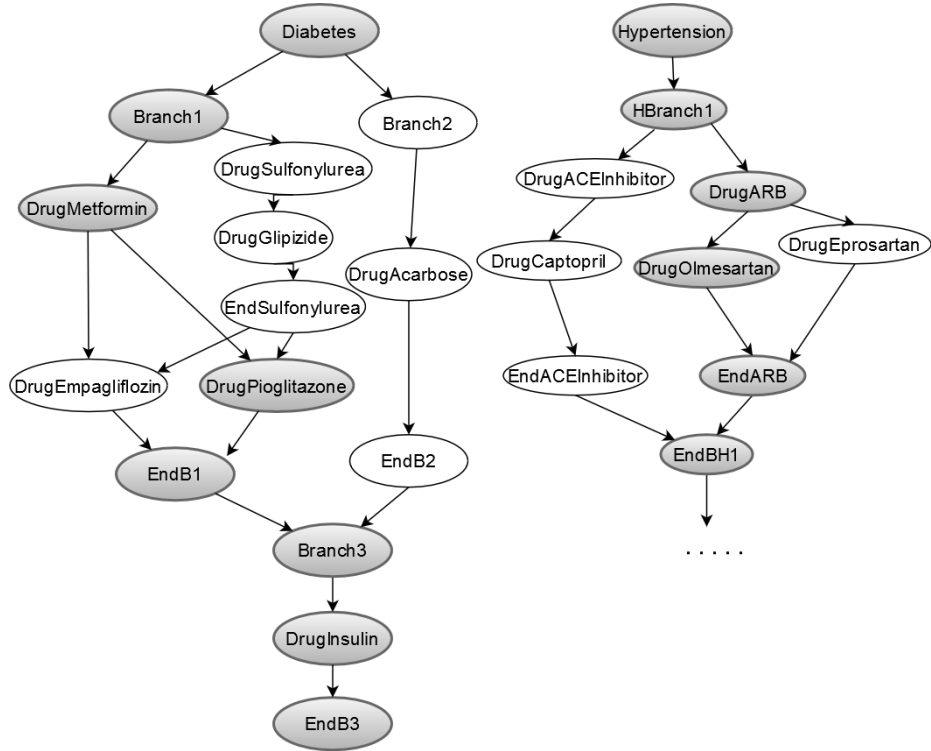


Fig. 1. PGraphs for Diabetes and Hypertension.

sub-branch with all member drugs related semantically. A path on a PG maps directly to a PGraph after accounting for additional marking nodes. An example PGraph is provided in Fig. 1 where the grey shaded nodes represent a path.

In addition to different levels of explanations, we also require to encode text that can be presented to a user. We base the explanation texts on the drug information provided by NICE² and the intuitions given in NICE pathways themselves. We manually filter out relevant information from these sources and organize it into types and levels. Types group the explanations into categories such as justification for a treatment (IP), using graph structure as a justification (PGRAPH), drug-drug conflict (DD), drug-disease conflict (CD) and drug-patient (CP) conflict, while levels correspond to the abstraction layers in the PGraph. We code these reasons with symbols R_i .

2.3 Argumentation Framework

In explaining a PGraph, our system should be able to progressively explore the solution path along PGraph identified by the SMT solver, and provide reasons

² For details see <https://bnf.nice.org.uk/>

for a specific choice on request from the user. We also require the system to be able to provide reasons explaining why another path has not been taken. To this end, we have chosen to encode PGraphs as argumentation frameworks, another form of graph, where nodes represent arguments and edges represent conflicts between arguments, where an evaluation process allows the identification of alternative sets of conflict-free nodes. The requirement for our system is two-fold: the encoding of the PGraph with additional justifications in a structured logic-based framework from which we can extract arguments that justify the pathway chosen by the SMT solver; and the definition of an exploratory dialogue protocol to allow human-like engagement with the explanation model. In a separate line of work, we have formulated a structured argumentation framework that encodes the reasons directly into the framework and allows us to identify a justified solution for a treatment plan obtained via the SMT solver. In this paper, however, our focus is on the dialogue describing this justified solution, therefore we only present here an abstraction of the arguments that can be formulated about a PGraph. In order to achieve this, we make use of an abstract argumentation framework (AF) as proposed by Dung [11]. Such a framework is defined as $AF = \langle AR, att \rangle$, where AR is a set of arguments and att represents a binary relation on AR , i.e., $att \subseteq AR \times AR$.

We use the approach of reformulating preferred semantics as argument labellings by Caminada et al. [8] to generate discussions about acceptable as well as rejected arguments in explanatory dialogue games. Caminada et al. [8] propose the following definitions which are useful in this regard.

Definition 1. *For a framework $AF = \langle AR, att \rangle$:*

1. *A labelling is a total function, $\mathcal{L}ab : AR \rightarrow \{in, out, undec\}$.*
2. *$\mathcal{L}ab$ is an admissible labelling iff for each $A \in AR$:*
 - (a) *if $\mathcal{L}ab(A) = in$ then $\forall B \in AR : (B att A \supset \mathcal{L}ab(B) = out)$.*
 - (b) *if $\mathcal{L}ab(A) = out$ then $\exists B \in AR : (B att A \wedge \mathcal{L}ab(B) = in)$.*
3. *$\mathcal{L}ab$ is a preferred labelling iff it is an admissible labelling where $in(\mathcal{L}ab)$ and $out(\mathcal{L}ab)$ are maximal (w.r.t. set inclusion) among all admissible labellings.*
4. *A preferred extension is the set of in-labelled arguments of a preferred labelling.*

Here, we compute preferred extensions as above, and we then use the extensions as the basis of a dialogue game inspired by the approach of Shams et al. [22], where they use preferred semantics as the basis of a Socratic discussion. We map the optimal paths returned by the SMT solver to equivalent preferred extensions and use preferred labelling to generate discussion about in and out status of each argument representing a node on a specific path.

Internally we model three types of arguments in the AF: path arguments, explanation arguments, and prescription arguments. We assume that the conclusion of a **Path** argument is a node representing a drug or a graph structure marker such as a start branch node. An **Explain** argument for a drug is used to conclude that there is a justifications for prescribing the drug, and will include a specific reason R_i . A **Prescribe** argument concludes that if there is a

justified drug, the drug can be prescribed. **Path** arguments attack all other **Path** and corresponding **Prescribe** arguments that are not on the same path. A **Path** argument can also attack **Explain** arguments of nodes not in the selected path if the latter describes an ADR for that particular **Path** argument. Consequently, each extension includes arguments prescribing a particular drug on a path in the graph along with its corresponding explanations.

3 Dialogue Games for Explanations

Formally dialogue games are interactions between two or more players according to a pre-defined set of rules [17]. The rules describe commencement conditions, permitted utterances (locutions), permitted combinations of locutions, commitments of participants, and termination conditions for the game. Dialogue game protocols have been developed for most of the primary dialogue types identified by Walton and Krabbe [24] such as Information Seeking, Inquiry, Persuasion, Negotiation, Deliberation and Eristic Dialogue.

We present two explanatory dialogue game protocols based on the Information Seeking dialogue which involves one participant seeking an answer from the other participant, who the former believes knows the answer. The objective of the dialogue game is to justify why a node is included or not included in a path.

The game has two participants: an Oracle that disseminates information and a Seeker who is looking for explanations. The Oracle has the *AF* and the extension corresponding to the recommendation by the SMT solver as part of its knowledge base. The Oracle knows the preferred labelling and shares arguments related to the treatment pathway with the Seeker at the start of the game through a *claim*. This is a list of **Path** arguments (labelled $\mathcal{L}ab(Arg) = in$ nodes) along with the **Path** arguments labelled $\mathcal{L}ab(Arg) = out$; we refer to this set as \mathcal{S} where $(Arg, \mathcal{L}ab(Arg)) \in \mathcal{S}$ if $Arg \in AR$ and $\mathcal{L}ab$ is a preferred labelling.

Subsequently, the Oracle uses the explanations present in the extension to answer the questions by the Seeker. In our dialogue protocols, the Seeker corresponds to a Patient agent seeking explanations for recommended medications. We define two versions of the dialogue protocols, passive explanatory protocol and active explanatory protocol based on the role of the Seeker. In the former, the Seeker acts as a passive listener while in the latter, the Seeker plays an active role in the dialog by confirming or refuting the Oracle’s assumptions about its preferences. Passive protocol covers the essential interaction and can be used in a scenario where preference information is missing such as an agent seeking information on behalf of another. Active protocol, on the other hand, provides detailed interaction which takes preferences into account and models a real world medical practice of taking patient history. The subsequent sections formalise the protocols and provide example dialogues.

Each dialogue protocol has three stages: Commencement, Progress, and Termination. The protocol specifies what speech acts can be exchanged according to what was previously exchanged in the dialogue at each phase. In our model, this is represented with a table where the right-hand side column states what

speech act can be used in response to the speech act moved on the left-hand side column (see Table 1 and Table3). In the dialogue protocol, participants have a commitment store \mathcal{CS} which includes a set of arguments with which the participants have committed to stand by [17]. In our dialogue, the commitment store is formed by arguments exchanged and related attacks. During the dialogue, participants update the commitment store via an assertion $\text{Assert}(\text{Arg})$ such that the new commitment store result in $\mathcal{CS} = \mathcal{CS} \cup \text{Arg}$, assuming monotonic updates. If an argument Arg exchanged is a rejection of a previous argument, we assume that \mathcal{CS} is updated with a new attacking argument Arg' where $\text{Arg}' \text{ att } \text{Arg}$; this update is referred to as $\sim \text{Arg}$ for convenience. We provide example dialogues in Table 2 and Table 4 with the three dialogue stages marked. The tables show the locutions underlying each natural language sentence next to each sentence and highlight the corresponding explanation Types and Levels that the natural language explanation represents. The last column shows example assertions for the \mathcal{CS} of each participant including the assertion type and argument as defined in Table 1 and Table 3.

3.1 Passive Explanatory Dialogue Game

In this dialogue, the Seeker plays a passive role and only receives updates from the Oracle. Information flow is one sided. Fig. 2 shows an overview of the dialogue.

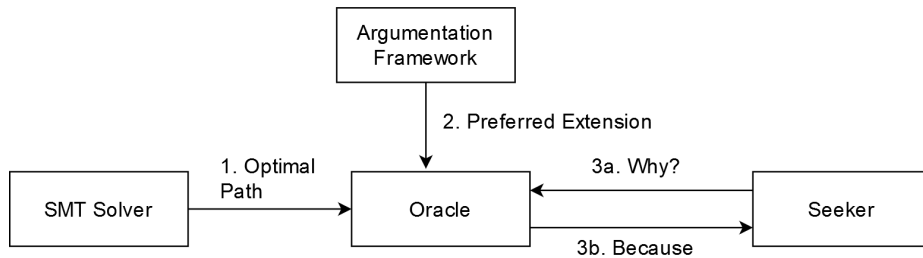


Fig. 2. Passive Explanatory Dialogue Overview.

Commencement The first move by Oracle (O) is a *Claim* which sets the context of the dialogue.

Progress

1. Every subsequent move by O is an explanation for the previous move of Seeker (P).
2. Each subsequent move by P is a query about an in or out label in the claim.
3. P can repeat any query until O forces it to move on by sending a *Move* locution.

Table 1. Passive Explanatory Dialogue Protocol Definition.

Request	Response
$Start_P$	$Claim_O(\mathcal{S})$
$Claim_O(\mathcal{S})$	$Why_P(query_i)$ where $query_i \in \mathcal{S}$ and query can be <i>in-query</i> or <i>out-query</i>
$Why_P(query)$	$Because_O(R_{ij})$ where R_{ij} represents all the reasons from 1 to i at abstraction level j . Different reasons at the same level can be combined using <i>and</i> or <i>alert</i> depending on the context. This response equals an assertion of one of the types defined as: $Assert(Arg)$: simple assertion where $Arg \in AR$. $Assert-Alt(Arg)$: asserting PGRAPH Type reason. $Assert-CD(Arg)$: asserting a drug-disease conflict as reason. $Assert-DD(Arg)$: asserting a drug-drug conflict as reason.
$Why_P(in-query)$	$Alert_O(R_{ij})$ where R_{ij} as previously explained. Alert message indicates a counter reason and represents a conflict that cannot be helped.
$But-Why_P(query)$	$Because_O(R_{ij-1})$ where $j-1$ is the next lower abstraction level than was used in previous $Because_O$ response for same query.
$But-Why_P(in-query)$	$Alert_O(R_{ij-1})$
$But-Why_P(query)$	$Move_O$ if $j-1 = 0$ where j equals the maximum number of levels defined.
$Move_O$	$Why_P(query_k)$ where $query_k \neq query_i$ for $query_k, query_i \in \mathcal{S}$ and $query_i$ has already been dealt with.
$Because_O$	End_P
$Alert_O$	End_P
End_P	No response. End dialogue session.

4. O provides all the explanations at the highest abstraction level in the graph first and moves on to the next lower abstraction level for subsequent repetitions of the query until it runs out of explanations. It then uses the Move locution.
5. O provides supporting reasons (IP Type) to justify an in status or counter reasons (CD, DD Types) to explain any possible conflicts that influenced the out label assignment.

Termination

1. P can terminate the dialogue at any time by using an *End* locution.
2. O can end the dialogue when all the nodes in the claim have been justified.

3.2 An example dialogue

We provide an example of a passive exploratory dialogue in Table 2. Subscripts O and P represent the Oracle and Seeker respectively and identify the participant making the move. Assume that the resulting argumentation framework for a path chosen by the SMT solver on the basis of a PGraph in Figure 1 is presented in Figure 3.

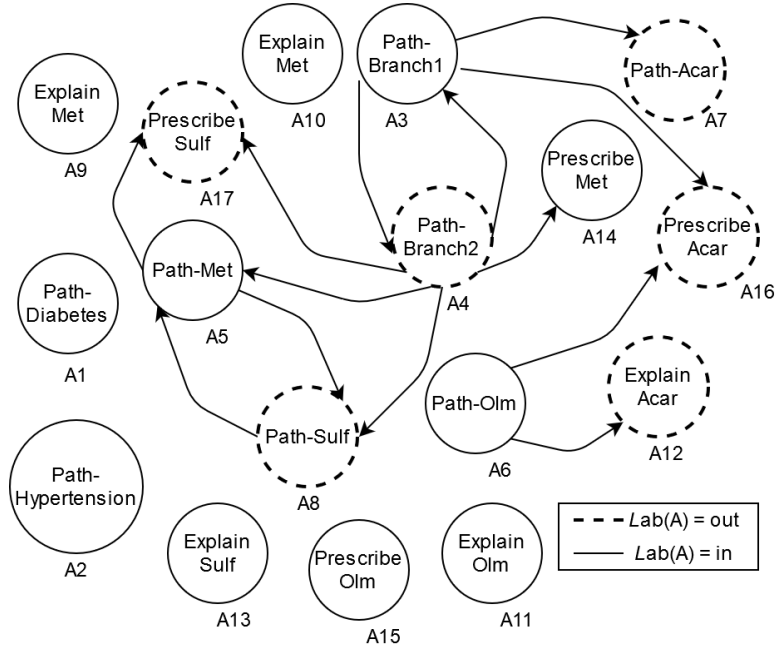


Fig. 3. Argument Graph for the AF of Table 2.

The AF is represented using an argument graph where the nodes represent arguments and the directed edges point towards the attacked nodes. Let arguments $A1..A8 \in AR$ represent edges in a PGraph, $A9..A13 \in AR$ represent corresponding justifications and $A14..A17 \in AR$ represent prescriptions for drugs represented by $A5..A8$. Arguments $\{A1, A2, A3, A5, A6, A9, A10, A11, A13, A14, A15\}$ represent a preferred extension which is part of the preferred labelling.

3.3 Properties of the Protocol

1. The protocol exploits the lack of path order in the extension by allowing P to query about any random node in the Claim.
2. The protocol preserves the graph abstraction levels qualitatively through abstractions of explanations.
3. O combines more than one explanations as a justification for the same query if multiple explanations exist at the same abstraction level.
4. The protocol ensures termination by preventing infinite loops using two strategies:
 - (a) O cannot repeat the same answer for the same query, rather it is forced to exhaust explanations at different abstraction levels progressively.
 - (b) P cannot repeat any query for which it has received a *Move* locution.
5. The protocol highlights any inconsistency between nodes in an extension through explanations by grouping supporting reasons for a node with counter reasons indicating possible adverse interactions.

Table 2. Passive Explanatory Dialogue example

Dialogue	Explanation Type/Level	Locutions	Commitment Store
Commencement:			
O: You have Diabetes and Hypertension. Let us go through your recommended NICE pathways for drug treatment. We will go through each drug one by one. Suggested solution is: S= [Diabetes, in(Metformin), out(Sulfonylurea), out(Acarbose), in(Olmesartan) ...]		Claim(S)	Assert(Diabetes) Assert(Hypertension) Assert(Metformin) Assert(Sulfonylurea) Assert(Acarbose) Assert(Olmesartan) ...
Progress:			
P: Why is Metformin suggested?		Why(in-Metformin)	
O: Because it is the default treatment on the recommended pathway	Type IP/Branch Level	Because	Assert(A8)
P: But why is Metformin suggested?		But-Why(in-Metformin)	
O: Because it is initial drug for treating Diabetes	Type IP/Drug Level	Because	Assert(A9)
P: Why is Sulfonylurea not suggested?		Why(out-Sulfonylurea)	
O: Because Metformin is an alternative to Sulfonylurea.	PGraph/ PGraph Level	Because	Assert-Alt (Metformin, Sulfonylurea)
P: But why is Acarbose not suggested?		But-Why(out-Acarbose)	
O: Because it adversely reacts with suggested drug Olmesartan.	Type DD/Drug Level	Because	Assert-DD(Acarbose, Olmesartan)
Termination:			
P: I have now received an update on the recommended pathway.End			

3.4 Active Explanatory Dialogue Game

In this case the Seeker actively participates in the conversation and confirms or refutes the assertions of the Oracle. The protocol follows a static update approach, where Seekers assertions do not take immediate effect. The Oracle passes the Seekers assertions to the SMT solver to get a more customized recommendation and uses that as the basis of a new discussion. In case Seekers assertions rule out all possible path options in the graph, the SMT solver returns the path that optimizes Seekers preferences even if it is not possible to meet all the requirements. Fig. 4 shows a schematic overview of the workflow.

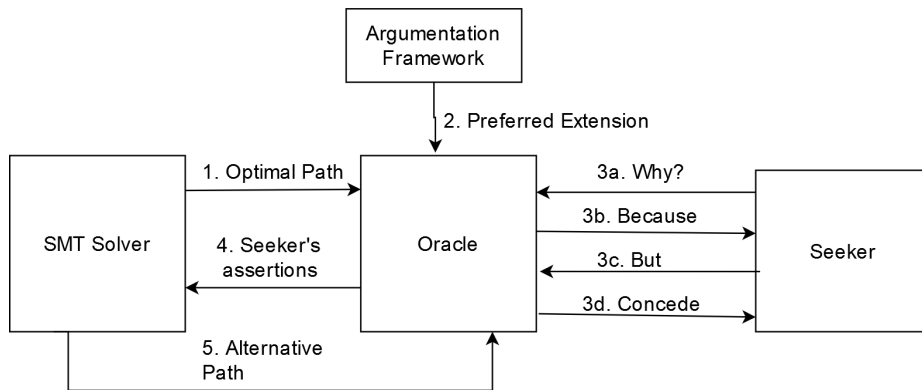


Fig. 4. Active Explanatory Dialogue Overview.

The dialogue starts off with the Oracle going through a checklist of items with the Seeker which capture Seeker preferences such as patient intolerance of specific drugs in a medical setting. We modelled the preferences for multimorbid patients by considering drug information available online and talking to experts. We learned that the presence of three conditions need to be considered before prescribing any drug: hepatic impairment, renal impairment, and pregnancy. This information is not part of the PGraph, rather the Oracle extracts this information from the Seeker at the start and passes it to the SMT solver which takes them into account when finding the suitable path.

Commencement

1. The first move by Oracle (O) marks the start of the *Checklist* phase in which each subsequent move by O is directed towards establishing preferences of the Seeker (P).
2. Each move by P in this phase confirms or refutes the assertion made by O in previous move.
3. O ends *Checklist* phase by send a *Claim* locution which proceeds as before.

Table 3. Active Explanatory Dialogue Protocol Definition.

Request	Response
<i>Checklist_O</i>	<i>Ready-Checklist_P</i>
<i>Assert-Item_O(Arg_i)</i>	<i>Assert-Item_P(Arg_i)</i> where $Arg_i \in AR$ is the i^{th} item on the check list and can equal $\sim Arg_i$ to show rejection of the assumption by the Seeker.
<i>Because_O(R_{ij})</i> with a simple assertion, <i>Assert_O(Arg)</i> as only this assertion type is defeasible. The rest represent facts which cannot be disputed by the Seeker.	<i>Assert_P(Arg)</i> where $Arg \in AR$. A positive Arg shows approval by Seeker whereas $\sim Arg$ represents rejection of the recommendation.
<i>Assert_P(Arg)</i>	<i>Concede_O</i>
<i>End_P</i>	<i>Alternative_O(S)</i>
<i>Alternative_O</i>	<i>Ready-Alternative_P</i>
<i>Alternative_O</i>	<i>Refuse_P</i>
<i>Accept_O(S)</i>	<i>Concede_P(S)</i>
<i>Ready-Alternative_P</i>	<i>Claim_O(S)</i>
<i>Accept_O(S)</i>	<i>Reject_P</i>
<i>Reject_P</i>	<i>Alternative_O</i>
<i>Refuse_P</i>	No response. End dialogue.

Progress

1. The progress rules for passive dialogue hold. Additionally each explanation for a query about an *in* labelled argument requires P to state its preference for the argument $Arg \in AR$ mentioned in the explanation by either confirming or refuting the assumption of O regarding Arg .
2. O can propose an alternative claim after all the labels in the first claim have been discussed. The dialogue then proceeds with the alternative claim as before.

Termination

1. P can end the dialogue at any time by using an *End* locution as before.
2. P accepts the proposed solution.
3. P refuses the offer of O to discuss an alternative claim.

Active Explanatory Dialogue Protocol extends Passive Dialogue Protocol with additional responses given in Table 3. Table 4 provides an example of how this dialogue starts.

3.5 Properties of the Protocol

1. O can use the assertions of P regarding its preferences as justifications when discussing the alternative claim.

Table 4. Active Explanatory Dialogue Game example.

Dialogue	Explanation Type/Level	Locution	Commitment Store
Commencement:			
O: Lets go through your basic health check list.		Checklist	
P: OK. I am ready		Ready	
O: Do you suffer from Hepatic Impairment?	Type CP/Drug Level	Assert-Item	Assert(C2)
P: No, I do not suffer from Hepatic Impairment.		Assert-Item	Assert(\sim C2)
O: Do you suffer from Renal Impairment?	Type CP/Drug Level	Assert-Item	Assert(C3)
P: No, I do not suffer from Renal Impairment.		Assert-Item	Assert(\sim C3)

2. Since the same query can be repeated, P must repeat the assertions. This could lead to an inconsistency if P changes the truth value for its assertion. This can be resolved in two ways:
 - (a) By preventing P from changing the truth value for the same query.
 - (b) By allowing P to change truth values and letting the last truth value hold. This last value can then be passed to the SMT solver.
3. O can use an estimate of how much P’s assertions contradict its world view and use a threshold value to proceed with finding an alternative claim. O could resolve minor objections by itself if the estimate is less than the threshold value.

4 Conclusion and Future Work

Explanations are an important tool to increase the trust we have on the recommendations of AI systems. Because of their critical importance, numerous approaches have been developed for clinical decision support such as expert systems [21], SMT solvers [5, 6], and argumentation [9, 15]. Here we use argumentation in an explanatory role to explain the decisions of the SMT solver.

We show how the recommendations of a black-box model like the SMT solver, can be justified through argumentation-based explanatory dialogue games. We describe how the dialogue games can provide explanations for the underlying reasoning of the SMT solver in calculating optimal paths which minimize ADRs. We also show that the games protocol is flexible and can highlight possible ADRs qualitatively. The approach allows engaging with the justifications behind the recommendations at various abstraction levels, resulting in a dialogue model which adapts to specific needs of the user. It also allows additional information

such as preferences to be introduced without overwhelming the user. The argument formulation is very expressive and allows both supporting reasons for selecting a node to be modelled as well as counter reasons expressed as a negative relationship between entities. The reasons themselves form part of the premise of the arguments, and are represented as propositions which are mapped to natural language representation during the dialogue. This approach can be used for interactive explanatory dialogues to explain similar sequence graphs in different domains when the underpinning justification model is provided.

In the future, we plan to divide the Oracle between multiple agent specialists as it will enable a more modular approach. Another possible direction is to explore the changes in the dialogue protocol if the Seeker's assertions have immediate effect. The current protocol takes the credibility of the agents for granted. It will be interesting to see how adding a credibility function can affect the dependability of the dialogue game. Most importantly, we plan to evaluate the approach through user studies involving health practitioners and patients alike and how they engage with our explanatory dialogue games. This will not only give us a mechanism to validate our approach, but also explore further arguments such as whether in some cases the optimal solution computed by an SMT solver is not in fact the preferred one.

References

1. Adadi, A., Berrada, M.: Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access* **6**, 52138–52160 (2018)
2. Biran, O., Cotton, C.: Explanation and justification in machine learning: A survey. In: *IJCAI-17 workshop on explainable AI (XAI)*. vol. 8 (2017)
3. Bjørner, N., Phan, A.D., Fleckenstein, L.: νz - an optimizing smt solver. In: Baier, C., Tinelli, C. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2015)*. LNCS, vol. 9035, pp. 194–199. Springer Berlin Heidelberg (2015)
4. Black, E., Atkinson, K.: Dialogues that account for different perspectives in collaborative argumentation. *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems* pp. 867–874 (2009)
5. Bowles, J., Caminati, M.B.: An integrated approach to a combinatorial optimisation problem. In: Ahrendt, W., Tapia Tarifa, S.L. (eds.) *Integrated Formal Methods*. LNCS, vol. 11918, pp. 284–302. Springer International Publishing, Cham (2019)
6. Bowles, J., Caminati, M., Cha, S., Mendoza, J.: A framework for automated conflict detection and resolution in medical guidelines. *Science of Computer Programming* **182**, 42 – 63 (2019). <https://doi.org/https://doi.org/10.1016/j.scico.2019.07.002>
7. Bowles, J.K.F., Caminati, M.B.: Balancing prescriptions with constraint solvers. In: Liò, P., Zuliani, P. (eds.) *Automated Reasoning for Systems Biology and Medicine, Computational Biology*, vol. 30, chap. 9, pp. 243–267. Springer, Cham (2019)
8. Caminada, M.W.A., Dvok, W., Vesic, S.: Preferred semantics as socratic discussion. *Journal of Logic and Computation* **26**(4), 1257–1292 (02 2014). <https://doi.org/10.1093/logcom/exu005>

9. Čyras, K., Oliveira, T.: Resolving Conflicts in Clinical Guidelines using Argumentation. Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS **3**, 1731–1739 (feb 2019)
10. De Vries, E., Lund, K., Baker, M.: Computer-mediated epistemic dialogue: Explanation and argumentation as vehicles for understanding scientific notions. The journal of the learning sciences **11**(1), 63–103 (2002)
11. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. Artificial Intelligence **77**(2), 321–357 (1995)
12. Fan, X., Toni, F.: On computing explanations in argumentation. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI’15). pp. 1496–1492. AAAI Press (2015)
13. Garca, A.J., Chesevar, C.I., Rotstein, N.D., Simari, G.R.: Formalizing dialectical explanation support for argument-based reasoning in knowledge-based systems. Expert Systems with Applications **40**(8), 3233 – 3247 (2013). <https://doi.org/https://doi.org/10.1016/j.eswa.2012.12.036>
14. Hughes, L.D., McMurdo, M.E.T., Guthrie, B.: Guidelines for people not for diseases: the challenges of applying UK clinical guidelines to people with multimorbidity. Age and Ageing **42**(1), 62–69 (08 2012)
15. Hunter, A., Williams, M.: Aggregating evidence about the positive and negative effects of treatments. Artificial Intelligence in Medicine **56**(3), 173–190 (2012). <https://doi.org/10.1016/j.artmed.2012.09.004>
16. Kovalov, A., Bowles, J.K.F.: Avoiding medication conflicts for patients with multimorbidities. In: Ábrahám, E., Huisman, M. (eds.) Integrated Formal Methods. LNCS, vol. 9681, pp. 376–390. Springer International Publishing, Cham (2016)
17. McBurney, P., Parsons, S.: Dialogue game protocols. In: Huget, M.P. (ed.) Communication in Multiagent Systems: Agent Communication Languages and Conversation Policies. LNCS, vol. 2650, pp. 269–283. Springer, Berlin, Heidelberg (2003)
18. de Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2008). LNCS, vol. 4963, pp. 337–340. Springer Berlin Heidelberg (2008)
19. de Moura, L., Bjørner, N.: Satisfiability modulo theories: Introduction and applications. Communications of the ACM **54**(9), 69–77 (sep 2011). <https://doi.org/10.1145/1995376.1995394>
20. Oliveira, T., Dauphin, J., Satoh, K., Tsumoto, S., Novais, P.: Argumentation with goals for clinical decision support in multimorbidity. In: Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS’18). pp. 2031–2033 (2018)
21. Rivas Echeverría, F., Rivas Echeverría, C.: Application of expert systems in medicine. In: Proceedings of the 2006 Conference on Artificial Intelligence Research and Development. p. 34. IOS Press, NLD (2006)
22. Shams, Z., De Vos, M., Oren, N., Padget, J.: Normative practical reasoning via argumentation and dialogue. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. pp. 1244–1250. AAAI Press (2016)
23. Tjoa, E., Guan, C.: A survey on explainable artificial intelligence (XAI): Towards medical XAI (2019)
24. Walton, D., Krabbe, E.: Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. State University of New York Press (1995)