



## King's Research Portal

*Document Version*  
Peer reviewed version

[Link to publication record in King's Research Portal](#)

*Citation for published version (APA):*

Sassoon, I. K., Kokciyan, N., Parsons, S. D., & Sklar, E. I. (2018). *Towards the use of Commitments in Multi-agent decision support systems*. Paper presented at The International Workshop on Dialogue, Explanation and Argumentation in Human-Agent Interaction (DEXAHAI), Southampton, United Kingdom.

### **Citing this paper**

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

### **General rights**

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

### **Take down policy**

If you believe that this document breaches copyright please contact [librarypure@kcl.ac.uk](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# Towards the use of Commitments in Multi-agent decision support systems

Isabel Sassoon  
Department of Informatics  
King's College London  
isabel.k.sassoon@kcl.ac.uk

Simon Parsons  
Department of Informatics  
King's College London  
simon.parsons@kcl.ac.uk

Nadin Kökciyan  
Department of Informatics  
King's College London  
nadin.kokciyan@kcl.ac.uk

Elizabeth I Sklar  
Department of Informatics  
King's College London  
elizabeth.sklar@kcl.ac.uk

## ABSTRACT

When a decision support system relies on information coming from different sources to reason about different recommendations it may be desirable to compute, record, track and engineer the influence that each source has on the elements which form part of the reasoning process. Different sources of information may have varying levels of reliability and, consequently, decision makers may want to reduce the potential impact of inputs from less reliable sources. In this paper, we introduce an approach for managing variable-quality inputs to a human multi-agent decision support system by using *computational argumentation* and *commitments* to model the relationships between sources and the data they provide, in a context-sensitive way. The methodology is illustrated with a clinical case study.

## CCS CONCEPTS

• **Computing methodologies** → **Multi-agent systems**; *Knowledge representation and reasoning*; *Distributed artificial intelligence*;  
• **Human-centered computing** → *Human computer interaction (HCI)*;

## KEYWORDS

Human Agent Interaction, Multi-Agent Decision Support, Computational Argumentation, Commitments

### ACM Reference Format:

Isabel Sassoon, Nadin Kökciyan, Simon Parsons, and Elizabeth I Sklar. 2018. Towards the use of Commitments in Multi-agent decision support systems. In *Proceedings of Workshop on Dialogue, Explanation and Argumentation at HAI 2018 (DEXA-HAI'18)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

DEXA-HAI'18, December 2018, Southampton, UK  
© 2018 Association for Computing Machinery.  
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00  
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

The main contributions that we articulate in this short paper are a method to apply commitments to an argumentation based decision support system in order to enable the influence of each different source of data to be tracked.

In some existing applications of Argumentation for Decision support [2], [5] arguments are instantiated by an agent using argumentation schemes that rely on data from different knowledge bases. In our scenario the data an agent requires to instantiate the relevant argument scheme is supplied by different agents. As such it is desirable that the provenance of the data supplied by the agents to instantiate an argument be recorded as part of the argument structure itself. We propose the use of commitments to cater for this requirement.

Decision-support systems do not always provide sufficient explanations to justify recommendations. In the context of health care this is crucial as health care professionals should have a complete picture of their patients in order to recommend an appropriate treatment. Hence, it is desirable to be able to trace back any recommendations made by a decision-support system to the evidence used to generate them and also the different sources of the evidence. In our case study, the *General Practitioner (GP)* will have seen an alert and an associated treatment recommendation for a stroke patient, Joey, based on all the data available (clinical history and sensor data). However, she may wish to remove the sensor data from this reasoning process as it is deemed unreliable. Consider the following case study:

Joey is a 63 year old male who after suffering a stroke is now monitoring his well-being using multiple wellness sensors. Joey experiences back pain and starts taking ibuprofen.

(t0) Joey notices that his blood pressure is elevated, based on the readings from his sensors, and books an appointment with his GP to discuss how to manage the elevated BP. He makes his readings available to the GP's dashboard.

(t1 - GP appointment) Given these readings the automated reasoning engine recommends modifying Joey's treatment by prescribing *angiotensin II receptor blocker (ARB)* to lower blood pressure.

(t2) Before modifying the treatment during the consultation the GP decides to take a few repeat ambulatory blood pressure readings, all of which are deemed to be normal and under the threshold for a hypertension diagnosis. Therefore, the GP does not trust the sensor

that Joey has been using, and requests that the decision support system re-evaluate recommendations in light of this. Hence, no new treatment is offered by the GP.

## 2 BACKGROUND

*Definition 2.1 (Argument).* An argument,  $Arg = \langle S, c \rangle$ , consists of a set of premises,  $S$ , defined in some language,  $\mathcal{L}$ , which support the conclusion,  $c$ .

An *Argument Scheme* is a model for instantiating arguments within a specific context and is used to provide a formal basis for instantiating arguments and defining their internal structure. An AS consists of a set of support premises ( $S$ ), which support the conclusion premise,  $c$ . necessary for this derivation. [10–12, 17]. Formally:

*Definition 2.2 (Argument Scheme).* An argument scheme  $AS = \langle S, c, \mathcal{V} \rangle$  consists of the set of premises,  $S$ , which support a conclusion,  $c$ , and are instantiated with the set of variables,  $\mathcal{V} = S.V \cup c.V$ .

An *Argumentation Framework (AF)* [4] represents a set of Arguments  $\mathcal{A}$ , and the relationships between the members of the set. Formally an *AF* is a pair  $\langle \mathcal{A}, \mathcal{R} \rangle$ , where  $\mathcal{A}$  is a set of arguments and  $\mathcal{R}$  is binary relation representing *attack* relationships between arguments. For example if  $Arg_1 = \langle p_1, c_1 \rangle$  and  $Arg_2 = \langle p_2, \neg c_1 \rangle$  then an attack relation exists between  $Arg_1$  and  $Arg_2$  since these arguments have conflicting conclusions (i.e. rebuttal attack).

## 3 METHOD

At the core of the decision support system are multiple agents, an agent reasoning about actions which relies on multiple other agents to supply the data required. In this clinical scenario we can assume the agents are: the treatment agent (instantiates the argument scheme for practical reasoning [1] as articulated in Table 1), a sensor agent (provides the data from a sensor), EHR agent (holds the patient facts from their electronic health record) and clinical guidelines agent (which given a specific patients facts holds the treatment guidelines). The Human interacting with this multi-agent system can be either the patient or the GP (or both).

For simplicity we will assume that the relevant clinical guidelines are extracted into a knowledge base (KB).

AS for practical reasoning
<i>premise</i> - In the current circumstances $R$
<i>premise</i> - We should perform Action $A$
<i>premise</i> - Which will result in new circumstances $S$
<i>premise</i> - Which will realise goal $G$
<i>premise</i> - Which will promote some value $V$
therefore : Action $A$ should be considered

**Table 1: Argument scheme practical reasoning from [1]**

An argument scheme can also be employed to map the inference process between findings (such as clinical measurements or observations) and a circumstance (such as a diagnosis). In our clinical scenario we are looking to infer whether high blood pressure is observed.

AS for defeasible rules
<i>premise</i> - if statements $P_1, P_2, \dots, P_n$ apply, then statement $Q$ may be inferred.
<i>premise</i> - Statements $P_1, P_2, \dots, P_n$ apply
therefore : $Q$ may be inferred

**Table 2: Argument scheme defeasible rules [3]**

AS for BP
<i>premise</i> - If mean blood pressure $M$ is higher than 140, High Blood Pressure can be inferred
<i>premise</i> - $M$ is higher than 140
therefore : High blood pressure ( <i>hbp</i> ) is inferred

**Table 3: Argument scheme for blood pressure measurements**

The *argument scheme for blood pressure measurements* (in Table 3) is a specialisation of the *argument scheme for defeasible rules (ASDS)* (in Table 2) introduced in [3] in the context of legal reasoning. The specialisation of the ASDS scheme involves replacing statements  $p_i$  with one statement regarding mean blood pressure measurements and  $q$  with a diagnosis.

The argument scheme for practical reasoning is associated with 16 critical questions, for the sake of simplicity we will focus on the first critical question only. *CQ1: Are the believed circumstances true?*

In the case study outlined in the first section, the current circumstances relate to elevated blood pressure readings from the sensors  $R$ , the action  $A$  is derived from a KB of treatments,  $S$  is normal blood pressure, and the goal  $G$  is to keep the patient healthy. The notion of value is not essential to this example and it has an overlap with the goal, so we will omit it.

The known facts about the patient are related to their elevated Blood Pressure:  $mean(BP)_{sensor1} > 140 \rightarrow R$ . The relevant action relates to prescribing a treatment  $A = ARB$  that will result in circumstances  $\neg R = S$ , and the goal  $G$  is Blood pressure control (*bpc*). Each of these relies on a different agent, which in turn relies on a different source of information. Respectively *sensor1*, *clinical guidelines* and *patient ehr*.

AS for practical reasoning
<i>premise</i> - In the current circumstances ( $bp\_mean(m), m > 140$ )
<i>premise</i> - We should perform Action <i>consider</i> ( $ARB$ )
<i>premise</i> - Which will result in new circumstances ( $bp\_mean(m), m < 140$ )
<i>premise</i> - Which will realise goal <i>BPC</i>
<i>premise</i> - Which will promote some value $V$
therefore : Action ( $ARB$ ) should be considered

**Table 4: Instantiated argument scheme practical reasoning from [1]**

Instantiating the Argument Scheme for practical reasoning as in Table 4 results in the argument:  
 $\langle \{ (bp\_mean(m), m > 140, action(arb), goal(bpc)) \}, consider(ARB) \rangle$ .  
 In this format we can see that there is no acknowledgement of the agent or source of the first premise (as an example).

$C_i$	debtor	creditor	antecedent	consequent	time
$C_1$	$Ag_{app}$	$Ag_{sensor}$	$\{bp\_mean(m), m > 140\}$	$observation(hbp) >$	$t_1-t_2$
$C_2$	$Ag_{app}$	$Ag_{gd}$	$\{observation(hbp), action(arb), goal(bpc)\}$	$consider(arb) >$	$t_2$

**Table 5: Instantiated commitments for the case study.**

### 3.1 Commitments

In our multi-agent decision support system, we need to model ways in which agents exchange information. A *commitment* [13] is a structure that can be used for tracing information exchanges over time. A commitment is an agreement that can be made between two agents, denoted as a 4-tuple:

$Commit(\text{debtor}, \text{creditor}, \text{antecedent}, \text{consequent})$

where the debtor agent makes a promise to the creditor agent to bring about the consequent, as long as the creditor agent brings about the antecedent.

The commitment instances for every interaction between two (or more) agents are saved in a *commitment store*,  $CS$ . Hence, agents can track which commitments are fulfilled or violated as a result of their interactions with other agents and who the other agents are. Note that this usage of *commitment store* aligns with prior use of the same term in earlier work on *argumentation-based dialogue* [14], where agents essentially make verbal commitments to each other by participating in a dialogue.

Different from previous argumentation-based approaches, an agent does not provide an argument by only using its own knowledge base, but an agent also collects pieces of information from various agents, similar to works described in [8, 9]. To this extent, we believe that commitments help agents to add a provenance layer while reasoning with data.

**3.1.1 Instantiating Commitments.** We make use of ASs to construct commitments, where an agent can collect the premises of an argument scheme from other agents. Thus, an argument scheme can be developed and instantiated in a distributed way. We define the representation of a commitment using an argument scheme to hold the antecedent and consequent:

$C_{id}(\text{debtor}, \text{creditor}, AS.S, AS.c)$

where  $id$  is the commitment identifier,  $AS.S$  is the antecedent and  $AS.c$  is the consequent. When the commitment is instantiated, the variables in  $AS.V$  are assigned values from the language  $\mathcal{L}$  (as in Definition 2.1).

**3.1.2 Mapping to Argumentation Frameworks.** Once all the of the commitments relevant to a specific case are instantiated the attack relations are determined by identifying any conflicting conclusions. The commitments and such attack relations can be visualised in a similar manner to an argumentation framework.

### 3.2 Case study

In Table 5, we show the commitment instantiations for the case study. At time  $t_1$ , the agent  $Ag_{sensor}$  reports that the mean blood pressure of Joey is more than 140. This instantiates the commitment  $C_1$ , and the agent  $Ag_{app}$  adds a new fact,  $observation(hbp)$ , to its knowledge base. The agent also notifies Joey that he has high blood

$C_1 : (Ag_{app}, Ag_{sensor}, \{bp\_mean(m), m > 140\}, observation(hbp))$

$C_2 : (Ag_{app}, Ag_{gd}, \{observation(hbp), action(arb), goal(bpc)\}, consider(arb))$

**Figure 1: The arguments instantiated at time  $t_2$  for the case study represented as commitments**

pressure and he should see his GP. At time  $t_2$ , the GP uses the app to get some recommendations for Joey. The guideline agent  $Ag_{gd}$  suggests that  $arb$  should be considered as per the argument scheme for practical reasoning (Table 1). This instantiates  $C_2$ , and the agent  $Ag_{app}$  recommends the use of  $arb$  ( $consider(arb)$ ) in Joey's case. Figure 1 illustrates the commitments in an argumentation framework, there are no attacks in this framework to keep the case study simple.

The GP decides to take a few repeat ambulatory blood pressure readings and all of which are deemed to be normal and under the threshold for a hypertension diagnosis. She adds a new observation (*nonreliable(sensor)*) to the app. The agent  $Ag_{app}$  generates commitment arguments according to this new information. Given the sensor is deemed no longer reliable,  $CQ1$  prevents the agent from generating  $C_1$  argument. Since  $C_1$  is no longer valid,  $C_2$  cannot be generated. In other words, the existence of  $C_2$  was dependent on the existence of  $C_1$ . As a result, no new treatment is offered for Joey. As shown in this example, as the agent is aware of the pieces of information provided by other agents (i.e. creditor agents), it can easily remove the corresponding commitments from the commitment store. Note that all the commitments are stored in the commitment store that can be reached by the debtor and the creditor agents.

## 4 RELATED WORK

There are multiple examples of argumentation-based (specifically argument scheme-based) systems for clinical decision support. The DRAMA agent proposed in [2] is an argumentation-based approach to reasoning about patient treatment. This is similar to our setting as it deals with treatment recommendations and makes use of an argument scheme (AS) with its related critical questions, however it also relies on values to reason with audience specific value-based argumentation frameworks. Our approach differs from this in that it is multi-agent, and each agent source is tracked in order to be able to remove any unreliable sources from the reasoning process.

In *arguEIRA* [5], the objective was to use argumentation to flag anomalies in patient's reactions to treatments in the Intensive Care Unit. This system did not consider possible different sensors as sources to the argument schemes. In *Carrel+* [15], where the objective of the argumentation based tool was to supervise and validate

the deliberation process on organ transplant viability, the arguments are generated by humans based on schemes so no disparate sources were involved. In the approach proposed in [6] preferences are used as a method to express the relative strength or weakness of evidence underpinning arguments. The evidence sources are rated by clinicians and then exploited to reason and aggregate evidence.

More recently in [7] an argumentation based approach to supporting treatment recommendations based on sensor data, clinical guidelines and electronic health data records was proposed. The approach makes use of argument schemes to instantiate treatment options and extended argumentation frameworks to incorporate preferences in the reasoning process. The method proposed does assume that the findings observed and use as inputs to the argument schemes are all equally reliable. A multi-agent application to clinical decision support is also proposed in [18] where a knowledge model also leverages published clinical guidelines.

## 5 CONCLUSION AND FUTURE WORK

In this short paper we have outlined an approach to integrating the different sources of knowledge and data used in argumentation based decision support, through enhancing the structure of an argument by representing them as commitments. This ensures that every argument instantiated, includes the sources of data or knowledge that underlie its premises. We have made use of a very simple clinical scenario to highlight the aspects of the decision support process where this approach would add value. To the best of our knowledge this approach is novel.

In situations where there are multiple measurements (e.g. blood pressure, heart rate and activity) and multiple sensors are used (e.g. heart rate measured both using a *fitbit* and using a mobile phone) then each of these and their role in the commitment and subsequently any arguments would be tracked. This would enable measurements from each separate sensor to be tracked through the recommendation process. We plan to assess how our approach compares and complements that approaches such as the one in [16]. Our future work includes expanding on all of the required formal definitions and applying this to a more realistic and complex scenario, such as one where multiple sensors are used.

## ACKNOWLEDGMENTS

This work was funded by the UK Engineering & Physical Sciences Research Council (EPSRC) under grant #EP/P010105/1 CONSULT.

## REFERENCES

- [1] Katie Atkinson and Trevor Bench-Capon. 2007. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence* 171, 10-15 (2007), 855–874.
- [2] Katie Atkinson, Trevor Bench-Capon, and Sanjay Modgil. 2006. Argumentation for decision support. In *Database and expert systems applications*. Springer.
- [3] Trevor Bench-Capon and Henry Prakken. 2010. Using argument schemes for hypothetical reasoning in law. *Artificial Intelligence and Law* 18, 2 (2010), 153–174.
- [4] Phan Minh Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and  $n$ -Person Games. *Artificial Intelligence* 77, 2 (1995).
- [5] Maria Adela Grando, Laura Moss, Derek Sleeman, and John Kinsella. 2013. Argumentation-logic for creating and explaining medical hypotheses. *Artificial intelligence in medicine* 58, 1 (2013).
- [6] Anthony Hunter and Matthew Williams. 2010. Using clinical preferences in argumentation about evidence from clinical trials. In *Proceedings of the 1st ACM International Health Informatics Symposium*. ACM, 118–127.
- [7] Nadin Kökciyan, Isabel K. Sassoon, Anthony Peter Young, Martin Chapman, Talya Porat, Mark Ashworth, Vasa Curcin, Sanjay Modgil, Simon Parsons, and Elizabeth I Sklar. 2018. Towards an Argumentation System for Supporting Patients in Self-Managing their Chronic Conditions. In *AAAI Joint Workshop on Health Intelligence (W3PHIAI 2018)*.
- [8] Nadin Kökciyan, Nefise Yaglikci, and Pinar Yolum. 2017. An Argumentation Approach for Resolving Privacy Disputes in Online Social Networks. *ACM Transactions on Internet Technology (TOIT)* 17, 3, Article 27 (2017).
- [9] Nadin Kökciyan and Pinar Yolum. 2017. Context-Based Reasoning on Privacy in Internet of Things. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*.
- [10] Simon Parsons, Katie Atkinson, Zimi Li, Peter McBurney, Elizabeth I Sklar, Munindar Singh, Karen Haigh, Karl Levitt, and Jeff Rowe. 2014. Argument schemes for reasoning about trust. *Argument & Computation, Special Issue: Trust, Argumentation and Technology* 5, 2–3 (2014).
- [11] Simon Parsons, Peter McBurney, Elizabeth Sklar, and Michael Wooldridge. 2007. On the relevance of utterances in formal inter-agent dialogues. In *Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- [12] Simon Parsons, Michael Wooldridge, and Leile Amgoud. 2003. On the outcomes of formal inter-agent dialogues. In *2nd International Conference on Autonomous Agents and Multi-Agent Systems*, J. S. Rosenschein, M. Wooldridge, T. Sandholm, and M. Yokoo (Eds.). ACM Press, New York.
- [13] Munindar P Singh. 1999. An ontology for commitments in multiagent systems. *Artificial Intelligence and Law* 7, 1 (1999).
- [14] Elizabeth I Sklar and M Q Azhar. 2015. Argumentation-based Dialogue Games for Shared Control in Human-Robot Systems. *Journal of Human-Robot Interaction* 4, 3 (2015). <https://doi.org/10.5898/JHRI.4.3.Sklar>
- [15] Pancho Tolchinsky, Ulises Cortes, Sanjay Modgil, Francisco Caballero, and Antonio Lopez-Navidad. 2006. Increasing human-organ transplant availability: Argumentation-based agent deliberation. *IEEE Intelligent Systems* 21, 6 (2006).
- [16] Alice Toniolo, Federico Cerutti, Nir Oren, Timothy J Norman, and Katia Sycara. 2014. Making informed decisions with provenance and argumentation schemes. In *Proceedings of the Eleventh International Workshop on Argumentation in Multi-Agent Systems*.
- [17] Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- [18] Liang Xiao and John Fox. 2017. A knowledge model and its application to agent-oriented clinical decision support. In *Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), 2017 10th International Congress on*. IEEE, 1–8.