A Collaborative Decision Support Tool for Managing Chronic Conditions

**Nadin Kökciyan[a], Martin Chapman[b], Panagiotis Balatsoukas[b], Isabel Sassoon[a], Kai Essers[a],
Mark Ashworth[b], Vasa Curcin[a,b], Sanjay Modgil[a], Simon Parsons[a], Elizabeth I Sklar[a]**

[a] *Department of Informatics, King's College London, London, United Kingdom,*
[b] *Department of Population Health Sciences, King's College London, London, United Kingdom*

Abstract

This paper describes work to assess the feasibility of using a decision support tool to help patients with chronic conditions, in particular stroke, manage their condition in collaboration with their carers and the health care professionals who are looking after them. The system contains several novel elements, in particular: the integration of data from commercial wellness sensors, electronic health records and clinical guidelines; the use of computational argumentation to track the source of data and to resolve conflicts and make recommendations; and argumentation-based dialogue to support interaction with patients. The proposed approach is implemented as an application that can run on smart devices (e.g. tablets). The users have personalised dashboards where they can visualise their health data, and interact with a conversational chatbot providing further explanations about their overall well-being.

Introduction

The aim of the CONSULT (Collaborative mObile decisioN Support for managing mULtiple morbidiTies) project is to explore the feasibility of employing a collaborative decision-support tool to help patients suffering from chronic diseases to self-manage their treatment plans. By 'collaborative', we mean that the patient, carers and medical professionals work as a team to decide on the best treatment plan for the patient. To establish feasibility we are developing a system, called CONSULT, which connects a patient to wireless sensors that are gathering data about them, provides real-time updates of data from their *Electronic Health Record* (EHR) and can provide recommendations and explanations based on clinical guidelines. Separately CONSULT provides a connection for a patient's general practitioner (GP) to have access to information being gathered about the patient. Feasibility is being assessed both at a technical level, in terms of whether it is possible to construct a working system that connects these disparate elements together, and at a usability level, in terms of whether all the parties find the system to be helpful. We are not, at this stage, assessing clinical benefit.

The CONSULT system exhibits the following novel features: (1) integration of data from commercial wellness sensors, patient's EHR, inputs from *Health Care Professionals* (HCPs), and treatment guidelines to produce an adaptive care plan customised to the patient's current circumstances; (2) application of *computational argumentation* to structure and track the data from these disparate sources and identify reinforcing and conflicting information; and (3) interaction with patients via *argumentation-based dialogue* to ensure understanding of the information gathered in (1) and to address, and potentially resolve, any conflicts found in (2). The users have personalised dashboards where they can visualise their health data and interact with the system.

Methods

Motivation

The CONSULT project was motivated by evidence that engaging patients in the self-management of chronic conditions can be beneficial to their well-being [3; 9; 20]. Clinical colleagues suggested that a suitable target population for a study in self-management would be stroke survivors, with the aim of the study being secondary stroke prevention. This suggestion was supported by an initial focus group with patients/carers and HCPs. In this focus group, stroke survivors reported a desire to receive additional support, beyond what can be provided by HCPs. In addition, HCPs at the focus group were keen to leverage new technology to help monitor patients.

The CONSULT System Overview

An overview of the CONSULT system architecture is shown in Figure 1. There are seven primary building blocks that make up the system: (a) patient input sources, including biometric data gathered by commercial wellness sensors and a patient's EHR; (b) user interfaces, including an interface for patients, as well as an interface for HCPs and a third interface for system administrators (orange blocks); (c) web-facing servers for gathering input data and supporting user interfaces (red blocks); (d) internal databases for storing raw data (blue blocks); (e) data mining processes, aggregating raw data and extracting natural language from arguments (yellow blocks); (f) aggregated data, including the output of the data mining and argumentation processes (pink blocks) and (g) a computational argumentation engine and associated sub-components, including inputs of computational guidelines, and drug interactions (green blocks). In the following sections, we describe the multiple information sources shown in Figure 1 in more detail, before detailing how these information sources are combined for the purpose of decision support. We then describe how a user can engage with the system.

Step (1.b) information provided by NICE guideline CG127:
"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."

Formal Representation in the Argumentation Engine:
$(age \geq 55) \lor black\text{-}african \lor black\text{-}caribbean \to offer(C, S_1, d)$
$\neg tolerated(C) \to \neg offer(C, S_1, d) \land offer(D, S_1, d)$
$oedema \lor heart\text{-}failure \lor highrisk\text{-}heart\text{-}failure \to offer(D, S_1, d)$

Figure 2 - Example Representation of Step 1 of NICE Guideline CG127

Information Sources

Guidelines

Clinical guidelines are documents that help HCPs and patients to decide on appropriate treatments. However, guidelines are mostly expressed in natural language. Clinical guidelines should be represented in a structured way in order to automate the reasoning process in decision-support systems. We represent domain-specific knowledge (e.g. the hypertension domain) using a logical language. We also use existing semantic representations of guideline information (e.g. drug interactions [26]) in the reasoning process.

Figure 1 - Abstract CONSULT System Architecture

Domain-Specific Representation: In order for CONSULT to reason about treatment plans, we represent knowledge in the hypertension domain using first order logic [12]. For example, in Figure 2, we represent part of the hypertension treatment guideline CG127 published by NICE [19]. The information provided for this particular step is represented in terms of logical rules. Patient characteristics, such as ethnicity or experienced side effects, could change the treatment plan. Hence, we take a similar approach to represent the relations between possible treatment plans and side-effects formally.

Use of External Ontologies: Patients typically deal with multiple morbidities, which makes the reasoning process difficult as conflicts among recommendations may arise. This requires the representation of: (1) recommendations that can be made for each condition (as described previously) and (2) potential interactions among such recommendations. Zamborlini et al. [26] introduce a semantic approach to detect interactions among recommendations by combining multiple guidelines. For CONSULT, we packaged Zamborlini's work as a web service, allowing us to create additional computational forms of guidelines in the semantic format

required for the identification of interactions. This information is then used as an additional data source—*Drug Interaction Finder* in Figure 1—for the argumentation engine. Specifically, new guidelines are authored as *quad triples*, added to a triplestore, and then processed by a logic-based reasoner in order to identify interactions of interest to the argumentation engine. The web service allows the argumentation engine to interrogate various stages of the interaction identification reasoning process, such as which recommended drugs have been identified as being in conflict.

Electronic Health Records

CONSULT's next information source is a patient's EHR specifically their demographic information, blood pressure history, medication history and details of long term conditions. To integrate with a given patient's EHR, we rely on an endpoint being provided by the vendor responsible for storing that patient's record. Typically, this endpoint is a local application programming interface (API) provided by each individual installation of the vendor's EHR software within a GP clinic. Therefore, CONSULT leverages the *data node connector* approach proposed in the TRANSFoRm project [5], by installing software within each GP practice that is designed to access this local API and transmit extracted EHR data to the rest of the CONSULT system for reasoning. Other models of EHR data access leveraged include direct collection of data on multiple patients from an external server provided by the vendor and simulated local API access using the N3/HSCN network. In any route to access the data, issues of governance are handled directly with the GP practice and the patient.

As each EHR vendor uses different formats to structure and code their data, once EHR data has been collected for a given patient, it needs to be standardized in order to enable the rest of the CONSULT system to be agnostic to the vendors from which the EHR data is derived. We choose Fast Healthcare Interoperability Resources (FHIR) standard as this format [2], and structurally transform each EHR to FHIR using a semi-automated matching and mapping process, while relying on services such as METMAPS [13] for code transformations, specifically to SNOMED which is used as part of the FHIR standard. Once transformed, data is inserted into a FHIR server, enabling the CONSULT system to operate as an *application* under the SMART-ON-FHIR paradigm [14].

Wireless Sensors

Our final data source is a patient's biometric health measures, extracted from a range of wearable devices. We primarily aim to acquire data on a patient's current blood pressure, pulse rate, activity and heart rate, since these are the most important measures for stroke patients. However, we do not ignore additional data that is also sent by the devices alongside this primary data (e.g. sleep quality). We employ a range of devices for this purpose, ranging from devices that are readily accessible to consumers (e.g. wrist worn devices) to more specialist medical devices (e.g. dedicated heart rate
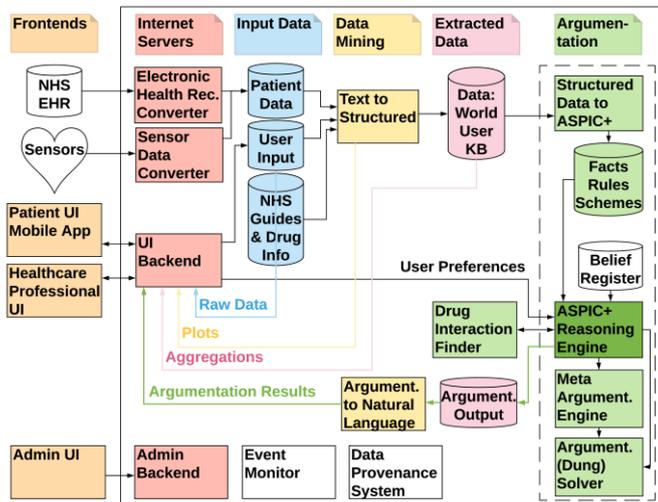
monitoring devices), where the former is advantageous as it increases the accessibility of the system, and the latter potentially offers greater accuracy and frequency of readings. A separate study involving some of the authors is looking to evaluate the quality of the sensor data produced by consumer medical devices.

In general, we aim to make the remainder of the CONSULT system as agnostic as possible to the hardware from which the readings originate. To do this, we build integration components for each wearable vendor's API—typically a remote REST endpoint, or a simpler data store—and then, upon the receipt of new sensor readings, convert this data from its vendor specific format to FHIR, which is also designed to represent and store patient health measures. This information can then be accessed by the rest of the system in the same manner as the EHR data via our FHIR server.

Integrating Data for Decision Support

The different patient data sources available to CONSULT (see Figure 1) are exploited and combined to present an up to date view of the patient's situation, as well as to support any reasoning in support of recommendations made. The data is merged and transformed in order to be able to monitor how the patient's latest readings compare to what is normal for the patient. In cases where there is a deviation, then relevant alerts are pushed to the GP and the patient accordingly.

The argumentation engine in CONSULT is the component where recommendations are made. This engine generates possible arguments and conflicts between them (e.g. conflicts in treatment guidelines that frequently arise in the management of multiple morbidities), and it computes treatment options to follow by providing further explanations for each of the options. For this, we make use of argument schemes [24] and critical questions to automatically construct arguments and identify conflicts between them. Argument schemes are semi-formal representations of the structures of common types of arguments, and they are useful in explaining why a particular argument was constructed. The argumentation reasoning engine, which is based on ASPIC+ [17], uses the received data to instantiate argument schemes and attack schemes in a metalevel argumentation framework [16; 25], and it constructs arguments and attacks to support any self management or treatment query related to the patient [12]. Such queries are submitted through the personalised dashboards of CONSULT, where the argumentation results are shared in a human-understandable way, and the stakeholders can interact with CONSULT to understand the decisions made by the argumentation engine.

User Interface and Interaction Scenarios

The interface for the CONSULT system has two main components: (1) a dashboard component that visualises longitudinal personal health data, presents tailored health recommendations to patients for disease self-management, and communicates the effect of different treatment and preventive interventions on their health risk (e.g. the risk of experiencing another stroke); (2) a conversational agent (chatbot) component the role of which is to provide patients with alerts and explanations about their health state (e.g. an increase in systolic blood pressure beyond the ideal reference range), to present treatment recommendations for self-managing their condition (e.g. which over the counter painkiller is the most indicative of reducing their backache given their current blood pressure levels, treatment plan and clinical guidelines), or, to allow users to perform, in an interactive environment, simple health information seeking tasks (e.g. in the form of acquiring

links to authoritative health literature and websites about a specific medication, measurement or condition).

In the remainder of the paper we will refer to the following example of Martin, a 60 year old male who has suffered a stroke, and who is using the CONSULT system and a variety of wellness sensors to monitor his own health. Martin and his GP interact through the CONSULT system.
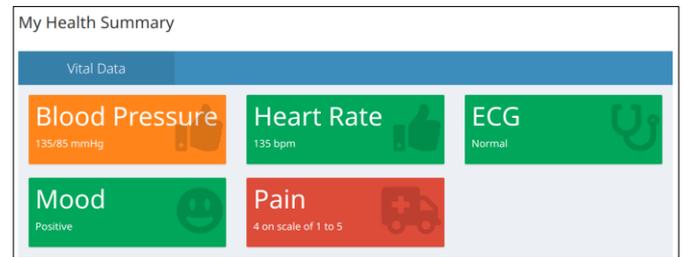

Figure 1 - The Dashboard (Overview) for an Android Tablet

Data Summary

The dashboard component of CONSULT contains an overview and a preview interface personalised for patients/carers and HCPs [4]. The overview interface, as depicted in Figure 3, displays a summary of the most recent measurements for all types of personal health data collected from the patient (e.g. blood pressure, heart rate, sleep activity, pain, stress, mood). For the representation of this information in the dashboard, we use a tile-based design where each tile provides information about a specific type of health data. Moreover, we use colour-coding to make clear immediately to the user when a specific measurement is outside the normal range [27]. For example, in the case of blood pressure, the colour green was used to indicate that the latest measurement was within the specified normal range, the colour orange indicated pre-hypertension levels—as depicted in Figure 3—, while red required attention. By selecting a tile from the overview interface the user can access longitudinal health data about the specific measurement at the preview interface. A typical preview interface provides users the opportunity to view their data at specific time intervals (e.g. hourly, daily, weekly, monthly or yearly), as averages or all raw measurements (using line graphs in the case of averages and scatter plots in the case of raw measurements). Also, for each time interval, the user is provided with additional descriptive information, such as the average, minimum and maximum value. In addition to personal health data, the dashboard provides users with the opportunity to use a risk calculator and to visualise (using *cates* plots) the effect of specific treatment and life-style interventions on their current risk of experiencing another stroke [8].

Actions were taken to improve the legibility and readability of content. These were: clean typeface (Arial), large default font size (12<), high contrast between characters and background (plain background and use of balanced colour saturation and luminance for text and graphs), writing that corresponds to a US 6th grade reading level (equal to year 7 for England). In terms of accessibility, actions were taken to improve access to content for colour blind users, including the use of both colour and symbols/labels to show that a value is within or outside normal range, or the selection of few well contrasting colours as opposed to multiple colours.

Treatment Recommender

If Martin's blood pressure is not under control, then as part of the consultation with his GP there needs to be a decision as to how to modify his treatment. The CONSULT system can

support this by presenting the GP with relevant, summarised and up to date patient data, and by presenting recommendations for possible treatments that consider this data, the patient's EHR, clinical guidelines and preferences of the stakeholders. The treatment recommendations are generated through the argumentation reasoning engine. A more in-depth description of the approach CONSULT takes in order to reason with the different possible and at time conflicting treatment options is in [12].

Interacting with the CONSULT ChatBot

The conversational component of the CONSULT system serves two main purposes. The first purpose is to provide a patient with the opportunity to seek immediate evidence-based advice about a specific health problem. For example, Martin may be suffering from back pain and CONSULT, using a chatbot, can advise him as to what he can do, as depicted in Figure 4. The chatbot is aware of the patient's latest wellness sensor readings, the data in their EHR—so will not recommend a treatment that has in the past caused side effects, for example—and clinical recommendations. These interactions are supported by argumentation-based dialogue [15; 22]. Additionally, the patient may have questions regarding their current treatment plan (e.g. why a particular medication has been prescribed). All the explanations are generated by the argumentation engine and displayed on the personalised dashboard. The second purpose of the conversational component is to alert the patient to an irregularity in one or more of their recent measurements and initiate a conversation, the purpose of which is to find a possible solution—suggesting the patient to review her blood pressure readings—or to advise the patient to contact a HCP.

Figure 2 - The interaction between Martin and the Chatbot

Results

We are currently in the process of evaluating the CONSULT system. Our intention was to design an application broad enough to accommodate the needs of people suffering from different chronic conditions, we have been focused on the context of stroke patients [20]. Based on our previous experience with this group of patients and the strong links to the South London Stroke Register (SLSR), we identified patients with different characteristics in terms of risk factors, comorbidity or demographic groups. The focus groups also involved co-design activities, following a design thinking approach [6], that resulted in user-generated versions of how

information should be displayed. We will conduct additional user studies to evaluate the usability of the proposed system to answer the following questions: (1) do the stakeholders of the CONSULT system (patients and HCPs) use the system? (2) do they think that it is useful to assist them in making decisions? and (3) do they like interacting with the system through the chatbot? Our initial focus groups have already allowed us to explore the answers to some of these questions [20].

Additionally, in [12], we have shown that argumentation is promising in explaining decisions to help HCPs and patients choose a treatment plan together. The use of argument and attack schemes specialised for the medical domain will be a next step to consider to generate better explanations [25]. The CONSULT system collects data from multiple information sources; as such, it is important to represent the interactions between these sources. One way of doing this is the use of commitments, which help the system to automatically decide what information source to trust and reason accordingly [21].

Discussion

Several works combine patient and clinical data collected from a variety of sources for the purposes of decision support, however many do not consider the number and variety of sources that are integrated by the CONSULT system. Systems that use a subset of the sources found in CONSULT include those that rely predominantly on sensor data, such as the system proposed by Groat et al. [11], which integrates data collected from glucose and exercise monitors to determine if patients are adhering to self-reported self-management behaviours. Others rely predominantly on a patient's medical history, such as the system proposed by Evans et al. [7], which aims to identify COPD in patients through a range of offline sources, including EHR data and echocardiograms, and the system proposed by Mosa et al. [18], which aims to identify patients at risk of CINV by mining EHR data.

In respect of reasoning with data sources, various works focus on developing argumentation-based systems for clinical decision support. Atkinson et al. propose the DRAMA agent to reason about patient treatment [1]. This is similar to our setting as it deals with treatment recommendations and makes use of argument schemes to construct arguments; however, each argument is associated with a value and the argumentation results change according to the prioritisation of such values. In *arguEIRA* [10], the authors make use of argumentation to detect and label anomalies in patient's reactions to treatments in the Intensive Care Unit. In *Carrel+* [23], the goal is to develop an argumentation based tool where agents conduct a deliberation dialogue to decide on the organ transplant viability. In contrast to these works, we consider data coming from multiple information sources rather than a centralised database. CONSULT also goes further than previous work in the degree to which it allows stakeholders to interact with the system to understand the argumentation reasoning better. CONSULT also provides dashboards to help patients self-manage their conditions and so provides a health monitoring facility that goes beyond the previous work cited.

Conclusions

CONSULT is one of the few systems to take a collaborative approach to the management of chronic disease. It is also the first decision-support system to make recommendations by combining multiple information sources, data science techniques, agreement technologies and an interactive chatbot.

We implement our proposed approach as a mobile application for Android tablets to help stroke patients and HCPs make decisions during the treatment process.

Future work focuses on the full evaluation of the system as a feasibility study for the deployment of this kind of technology. The main questions that need to be resolved are the technical feasibility of successfully operating a system that connects patient, sensors, EHR and GP together in real-time, and the feasibility of having patients, carers and medical professionals use the system without finding it burdensome. We believe that the principles behind CONSULT can be adapted to help with a number of chronic diseases, and we hope to explore this hypothesis in future work.

Acknowledgements

References

[1] K. Atkinson, T. Bench-Capon, and S. Modgil, Argumentation for Decision Support, in: *Database and Expert Systems Applications*, 2006, pp. 822-831.

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

[3] C.M. Boyd and M. Fortin, Future of multimorbidity research: how should understanding of multimorbidity inform health system design?, *Public Health Reviews* **32** (2010), 451.

[4] B. Brown, P. Balatsoukas, R. Williams, M. Sperrin, and I. Buchan, Interface design recommendations for computerised clinical audit and feedback: hybrid usability evidence from a research-led system, *International Journal of Medical Informatics* **94** (2016), 191-206.

[5] B.C. Delaney, V. Curcin, A. Andreasson, T.N. Arvanitis, H. Bastiaens, D. Corrigan, J.F. Ethier, O. Kostopoulou, W. Kuchinke, M. McGilchrist, P. van Royen, and P. Wagner, Translational Medicine and Patient Safety in Europe: TRANSFoRm-Architecture for the Learning Health System in Europe, *BioMed Research International* **2015** (2015), 1-8.

[6] K. Dorst, The core of 'design thinking' and its application, *Design studies* **32** (2011), 521-532.

[7] R.S. Evans, J.F. Lloyd, V. Flint, B.D. Horne, S. Rea, D.S. Collingridge, S.E. Abplanalp, A.A. Fazili, and D.P. Blagev, Decision Support to Help Identify Patients with Chronic Obstructive Pulmonay Disease Exacerbation, in: *American Medical Informatics Association (AMIA) Annual Symposium*, 2018, pp. 1687-1689.

[8] A. Fagerlin, B.J. Zikmund-Fisher, and P.A. Ubel, Helping patients decide: ten steps to better risk communication, *Journal of the National Cancer Institute* **103** (2011), 1436-1443.

[9] P. Fraccaro, M.A. Casteleiro, J. Ainsworth, and I. Buchan, Adoption of clinical decision support in multimorbidity: a systematic review, *JMIR medical informatics* **3** (2015).

[10] M.A. Grando, L. Moss, D. Sleeman, and J. Kinsella, Argumentation-logic for creating and explaining medical hypotheses, *AI in Medicine* **58** (2013), 1-13.

[11] D. Groat, H. Soni, M.A. Grando, and D.R. Kaufman, Design and Testing of a Smartphone Application for Real-Time Self-Tracking Diabetes Self-Management Behaviors, *Applied Clinical Informatics* **9** (2018), 440-449.

[12] N. Kökciyan, I. Sassoon, A.P. Young, M. Chapman, T. Porat, M. Ashworth, V. Curcin, S. Modgil, S. Parsons, and E. Sklar, Towards an Argumentation System for Supporting Patients in Self-Managing their Chronic Conditions, in: *Proceedings of the AAAI Joint Workshop on Health Intelligence (W3PHIAI 2018)*, 2018.

[13] S.F. Liang, J.F. Ethier, T. Porat, A. Tapuria, B.C. Delaney, and V. Curcin, MeTMapS-Medical terminology mapping system, *CEUR Workshop Proceedings* **1692** (2016).

[14] J.C. Mandel, D.A. Kreda, K.D. Mandl, I.S. Kohane, and R.B. Ramoni, SMART on FHIR: a standards-based, interoperable apps platform for electronic health records, *Journal of the American Medical Informatics Association* **23** (2016), 899-908.

[15] P. McBurney and S. Parsons, Dialogue Games for Agent Argumentation, in: *Argumentation in Artificial Intelligence,* G. Simari and I. Rahwan, eds., Springer US, Boston, MA, 2009, pp. 261-280.

[16] S. Modgil and T. Bench-Capon, Metalevel Argumentation, *Journal of Logic and Computation* **21** (2011), 959-1003.

[17] S. Modgil and H. Prakken, A general account of argumentation with preferences, *Artificial Intelligence* **195** (2013), 361-397.

[18] A.S.M. Mosa, A.M. Hossain, and I. Yoo, A Dynamic Decision Support System for Preventing Chemotherapy-Induced Nausea and Vomiting, in: *American Medical Informatics Association (AMIA) Annual Symposium*, 2018.

[19] National Institute for Health and Care Excellence (NICE), Hypertension in adults: diagnosis and management, in: *https://www.nice.org.uk/guidance/cg127*, [accessed on March 28, 2019], 2011.

[20] T. Porat, N. Kökciyan, I. Sassoon, A.P. Young, M. Chapman, M. Ashworth, S. Modgil, S. Parsons, E. Sklar, and V. Curcin, Stakeholders' views on a collaborative decision support system to promote multimorbidity self-management: barriers, facilitators and design implications, in: *American Medical Informatics Association Annual Symposium*, 2018.

[21] I. Sassoon, N. Kökciyan, S. Parsons, and E.I. Sklar, Towards the Use of Commitments in Multi-agent Decision Support Systems, in: *Workshop on Dialogue, Explanation and Argumentation at Human Agent Interaction (HAI)*, 2018.

[22] E. Sklar and M.Q. Azhar, Argumentation-based Dialogue Games for Shared Control in Human-Robot Systems, *Journal of Human-Robot Interaction* **4** (2015), 120-148.

[23] P. Tolchinsky, U. Cortes, S. Modgil, F. Caballero, and A. Lopez-Navidad, Increasing human-organ transplant availability: Argumentation-based agent deliberation, *IEEE Intelligent Systems* **21** (2006), 30-37.

[24] D. Walton, C. Reed, and F. Macagno, *Argumentation Schemes*, Cambridge University Press, 2008.

[25] A.P. Young, N. Kökciyan, I. Sassoon, S. Modgil, and S. Parsons, Instantiating Metalevel Argumentation Frameworks, in: *Proceedings of the 7th International Conference on Computational Models of Argument* 2018, pp. 97-108.

[26] V. Zamborlini, M.d. Silveira, C. Pruski, A.t. Teije, E. Geleijn, M.v.d. Leeden, M. Stuiver, and F.v. Harmelen, Analyzing interactions on combining multiple clinical guidelines, *AI in Medicine* **81** (2017), 78-93.

[27] B.J. Zikmund-Fisher, A.M. Scherer, H.O. Witteman, J.B. Solomon, N.L. Exe, B.A. Tarini, and A. Fagerlin, Graphics help patients distinguish between urgent and non-urgent deviations in laboratory test results, *Journal of the American Medical Informatics Association* **24** (2016), 520-528.

Address for correspondence

Simon Parsons, Department of Informatics, King's College London, Bush House, Strand Campus, 30, Aldwych, London WC2B 4BG, simon.parsons@kcl.ac.uk.