Artificial Intelligence in Medicine 58 (2013) 1–13

Contents lists available at SciVerse ScienceDirect

Artificial Intelligence in Medicine

journal homepage: www.elsevier.com/locate/aiim

Argumentation-logic for creating and explaining medical hypotheses

Maria Adela Grando[a,*], Laura Moss[b,c,d], Derek Sleeman[b,c], John Kinsella[b]

[a] Division of Biomedical Informatics, School of Medicine, University California San Diego, 9500 Gilman Drive #0505, La Jolla, CA 92093-0505, USA
[b] Academic Unit of Anaesthesia, Pain, and Critical Care Medicine, School of Medicine, University of Glasgow, Wolfson Medical School Building, University Avenue, Glasgow G12 8QQ, Scotland, UK
[c] Department of Computing Science, University of Aberdeen, Meston Building, Meston Walk, Aberdeen AB24 3UE, Scotland, UK
[d] Department of Clinical Physics, University of Glasgow, Wolfson Medical School Building, University Avenue, Glasgow G12 8QQ, Scotland, UK

ARTICLE INFO

Article history:
Received 13 July 2012
Received in revised form
14 December 2012
Accepted 5 February 2013

Keywords:
Knowledge-based expert systems
Explanation
Argumentation logic
Ontology
Intensive care unit

ABSTRACT

Objective: While EIRA has proved to be successful in the detection of anomalous patient responses to treatments in the Intensive Care Unit, it could not describe to clinicians the rationales behind the anomalous detections. The aim of this paper is to address this problem.
Methods: Few attempts have been made in the past to build knowledge-based medical systems that possess both argumentation and explanation capabilities. Here we propose an approach based on Dung's seminal calculus of opposition.
Results: We have developed a new tool, arguEIRA, which is an extension of the existing EIRA system. In this paper we extend EIRA by providing it with an argumentation-based justification system that formalizes and communicates to the clinicians the reasons why a patient response is anomalous.
Conclusion: Our comparative evaluation of the EIRA system against the newly developed tool highlights the multiple benefits that the use of argumentation-logic can bring to the field of medical decision support and explanation.

Published by Elsevier B.V.

1. Introduction

EIRA (Explaining, Inferencing, and Reasoning about Anomalies) is an existing knowledge-based system [1] which identifies anomalies in a domain and generates potential explanations for these anomalies. To evaluate the approach, EIRA has previously been applied in the intensive care unit (ICU) domain to help users detect anomalous patient responses to treatment in the ICU and provides decision support for clinicians by generating explanations for the anomalies. In this application, EIRA was based on complex algorithms that reflect the problem solving methods used by ICU clinicians in detecting and resolving anomalies. While EIRA has proved to be very accurate [2], it lacks a justification system that could make explicit, in a user-friendly way, the complex reasoning behind the algorithms; such an extension would enhance the acceptance of the tool by domain experts and allow the system's knowledge and reasoning to be refined more easily; these are important weaknesses that need to be addressed. In [3] we replaced EIRA's algorithms to detect anomalies for a more flexible reasoning process based on logical argumentation while we kept the more complex EIRA algorithms that generated potential reasons for why the detected anomalies may have occurred. In this work we suggest an implementation process to completely replace EIRA's algorithms that generate possible explanations for why the detected anomalies may have occurred.

Argumentation allows reasoning with imperfect information by constructing and evaluating arguments relevant to alternative, and in some cases, conflicting, conclusions. While in standard logic we cannot infer from a knowledge base *p* and *not p*, in argumentation logic *p* and *not p* can co-exist. An important benefit of argumentation is that it makes the sources of inconsistency clearer, and also allows the course of an argument to be reported, so we can reason methodically in the face of conflict.

Medicine is a good example of a domain where beliefs are not always categorical; rules can have unknown or implicit conditions and consequently can be incomplete; conclusions can be contradictory; or more than one explanation for a medical hypothesis could be possible based on different medical studies. There seems to be two solutions to this sort of problem when designing systems. The first is to remove the inconsistencies and resolve the conflicts by a thorough belief revision. The second suggests that the inconsistency should not be eradicated but should rather be kept to provide insights into rational processes. In this paper we want to explore

* Corresponding author at: University of California San Diego, Department of Medicine, Division of Biomedical Informatics, 9500 Gilman Drive #0505, La Jolla, CA 92093-0505, USA. Tel.: +1 858 8227810; fax: +1 858 8227685.
E-mail addresses: mgrando@ucsd.edu, adelagrando@yahoo.com.ar (M.A. Grando), laura.moss@glasgow.ac.uk (L. Moss).

0933-3657/$ – see front matter. Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.artmed.2013.02.003

the benefits of this second view by proposing an enhanced version of EIRA where the process used to construct justifications is based on ontology-based argumentation reasoning.

While the benefits of the use of argumentation-based inference and decision making in medicine has been explained before [4,5], this work provides the opportunity to generate empirical evidence for those claims by comparing EIRA's algorithmic-based outcomes with the outcomes generated from its reimplementation based on argumentation logic. While this work is a continuation of [3] the main focus of this paper is to: (1) replace EIRA's more complex algorithms which generate potential explanations for anomalous patient responses to ICU treatments with an argumentation-based decision system; whereas in [3] we only replaced the much simpler algorithm to detect an anomaly, (2) analyze the advantages that the use of advanced argumentation-based computational capabilities have brought to the enhanced EIRA system, and (3) discuss the state of the art on argumentation-based tools to conjecture possible future extensions of EIRA.

The first stage of enhancing EIRA involved using Dung's seminal calculus of opposition to identify and express argument exchanges made by ICU clinicians. The obtained arguments and interactions (attacking, supporting arguments) were modelled using the argumentation service platform with integrated component (ASPIC) engine (http://aspic.acl.icnet.uk/, accessed 30.11.12).

The second stage involved replacing the original EIRA's algorithm to detect anomalies with an argumentation-based hypothesis generator module based on the schemes recognized during the previous stage. Here we revisit the description presented in [3] of the implementation of an argumentation system for detecting anomalies and we extend it by providing details of the additional work required for providing EIRA with an argumentation-based explanation system. The added ASPIC schemes were defined in terms of the ICU domain ontology used in the existing EIRA. The hypothesis generator allows inferring a list of possible hypothesis explaining a patient's anomalous response to ICU treatment, ordered by the strength of supportive medical evidence.

The last stage consisted of developing a natural language explanation system that converts the argumentation exchange generated by ASPIC into a natural language explanation.

To the best of our knowledge currently there is no decision support system, apart from the tool we explain here, which allows for the automatic, ontology-driven detection and explanation of anomalies. In the field of medicine, there are a few argumentation-based decision support systems which provide natural language explanations of the rationality of the decision processes, but as it is indicated in [4] those approaches are based on a more restricted argumentation system than Dung's calculus that we used for implementing our tool. Because these approaches can only partially capture the behaviour that is observed in natural patterns of arguments they are less expressive than the tool we propose here.

This paper is organized as follows: in Section 2 we explain the background to this work, introducing the EIRA system (Section 2.1) and related argumentation-based systems for the generation and explanation of medical hypotheses (Section 2.2), including ASPIC, the argumentation inference engine chosen for the implementation of arguEIRA. In Section 3 we present our results. In Section 3.1 we explain how we performed the mapping of arguments exchanged in the ICU context into the ASPIC tool. In Section 3.2 we summarize the way we implemented arguEIRA. Section 4 is devoted to discussion. In Section 4.1, we compare the outcomes of the purely algorithmic-based EIRA with arguEIRA to explain the advantages of the use of argumentation-based reasoning in the ICU context. In Section 4.2 we compare arguEIRA with similar argumentation-based tools. Finally, in Section 5, we present our conclusions and future plans to further extend arguEIRA.

2. Background

2.1. EIRA

Anomalies have long been thought to drive scientific discovery. Kuhn explained that "*discovery commences with the awareness of anomaly, i.e., with the recognition that nature has somehow violated the paradigm-induced expectations that govern normal science*" [6]. Thagard [7] and Darden [8] have indicated that the detection and eventual explanation (using domain knowledge) of such anomalies force scientists to revise their knowledge in a number of ways; from a minor refinement of a hypothesis to major changes of background scientific knowledge. Anomalies in the medical domain are of interest as they often point to the inadequacy of a currently held medical theory and suggest the refinement of the related theory; consequently this can provide the impetus for the discovery of further medical knowledge by clinicians. EIRA [9] has been applied in the medical domain to detect anomalous patient responses to treatment in the ICU domain and assists clinicians by providing potential explanations for why the anomaly may have occurred. The explanations generated by EIRA can be considered as an anomaly-driven refinement of the clinician's theory.

EIRA is believed to be a generic tool and comprises: a knowledge base consisting of several instantiated OWL ontologies (http://www.w3.org/2004/OWL/, accessed 30.11.12) in this case, describing the ICU domain and a Java program implementing strategies extracted from domain experts' protocols. For the initial application of EIRA, studies were performed in the intensive care unit domain to investigate the ways in which clinicians detected and explained anomalies [1]; these studies were then used as the basis for the strategies. As shown in Fig. 1, in this application, EIRA also has access to routinely recorded data containing physiological parameters, and drug and fluid infusions for each patient. When attempting to detect anomalies, EIRA identifies the drugs given to the patient at a particular time point from the patient's data and retrieves the anticipated effects of administering each drug from the ICU ontology. When the anticipated response(s) do not occur, the actual response observed in the data is noted (Fig. 2). For example, EIRA's anomaly detection algorithm may determine that:

Patient Data contains: Adrenaline is administered to a patient and this is followed by a decrease in the patient's mean arterial pressure (MAP).
EIRA's Knowledge Base contains: Adrenaline should increase MAP.
EIRA suggests that: A patient has responded anomalously to adrenaline.

To generate context-dependent hypotheses for why the detected anomalies may have occurred, EIRA proceeds with each of the implemented strategies (algorithms) and, if appropriate, explanations are presented to the user. For example, EIRA may provide the following explanation for the previous anomaly (by invoking the overall condition deterioration algorithm):

Patient Data contains: Adrenaline is administered to a patient and this is followed by a decrease in the patient's mean arterial pressure (MAP).
EIRA's Knowledge Base contains: Adrenaline should increase MAP.
EIRA suggests that: A patient has responded anomalously to adrenaline.