Artificial Intelligence in Medicine 55 (2012) 137–148

Contents lists available at SciVerse ScienceDirect

Artificial Intelligence in Medicine

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

Inconsistency as a diagnostic tool in a society of intelligent agents

Marjorie McShane[a,*], Stephen Beale[a], Sergei Nirenburg[a], Bruce Jarrell[b], George Fantry[b]

[a] Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD 21250, USA
[b] University of Maryland School of Medicine, Baltimore, MD 21201, USA

ARTICLE INFO

Article history:
Received 27 October 2011
Received in revised form 26 March 2012
Accepted 25 April 2012

Keywords:
Cognitive modeling
Intelligent agents
Virtual tutor
Clinical decision making support

ABSTRACT

Objective: To use the detection of clinically relevant inconsistencies to support the reasoning capabilities of intelligent agents acting as physicians and tutors in the realm of clinical medicine.
Methods: We are developing a cognitive architecture, OntoAgent, that supports the creation and deployment of intelligent agents capable of simulating human-like abilities. The agents, which have a simulated mind and, if applicable, a simulated body, are intended to operate as members of multi-agent teams featuring both artificial and human agents. The agent architecture and its underlying knowledge resources and processors are being developed in a sufficiently generic way to support a variety of applications.
Results: We show how several types of inconsistency can be detected and leveraged by intelligent agents in the setting of clinical medicine. The types of inconsistencies discussed include: test results not supporting the doctor's hypothesis; the results of a treatment trial not supporting a clinical diagnosis; and information reported by the patient not being consistent with observations. We show the opportunities afforded by detecting each inconsistency, such as rethinking a hypothesis, reevaluating evidence, and motivating or teaching a patient.
Conclusions: Inconsistency is not always the absence of the goal of consistency; rather, it can be a valuable trigger for further exploration in the realm of clinical medicine. The OntoAgent cognitive architecture, along with its extensive suite of knowledge resources an processors, is sufficient to support sophisticated agent functioning such as detecting clinically relevant inconsistencies and using them to benefit patient-centered medical training and practice.

1. Introduction

The term *inconsistency* tends to imply a negative evaluation of a state of affairs, the lack of a state of consistency that is the implied goal. However, in societies of people, and in societies of intelligent agents modeled as people, inconsistency is not always a detriment, it can actually serve as a diagnostic tool. For example, in the domain of clinical medicine, inconsistencies between test results and the doctor's hypothesis about what is wrong can suggest that the hypothesis was incorrect or that the test results were flawed, and inconsistencies between a doctor's observation and a patient's report can suggest an intentional or unintentional misrepresentation by the patient. Similarly, in the domain of *teaching* clinical medicine, inconsistencies between an expert's preferred approach to solving a problem and a novice's approach can suggest room for improvement for the novice. Accordingly, if intelligent agents are modeled to function as clinicians or as tutors for clinicians, they should be prepared to exploit *diagnostic inconsistencies* in the same way as people do. This paper describes our recent work in configuring intelligent agents who serve in both of these roles and count among their stockpile of cognitive capabilities the detection of diagnostic inconsistencies.

1.1. OntoAgent agents

The intelligent agent environment to be used for illustration is OntoAgent, which supports the modeling of human-like behavior in artificial intelligent agents that collaborate with people (http://www.trulysmartagents.org/index.php). OntoAgent is a multi-agent environment being developed to support a suite of applications, including training and advising in the domain of clinical medicine. The development approach, which seeks to model agents that function like people, necessitates that our program of R&D be comprehensive, actively covering areas including the following:

- Physiological simulation using hybrid domain knowledge: physiological pathways when they are known and relevant for the goals of the simulation, and clinical "bridges" when pathways are

* Corresponding author at: CSEE Dept., ITE 325, University of Maryland Baltimore County, Baltimore, MD 21250, USA.
E-mail address: marge@umbc.edu (M. McShane).

0933-3657/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.artmed.2012.04.005

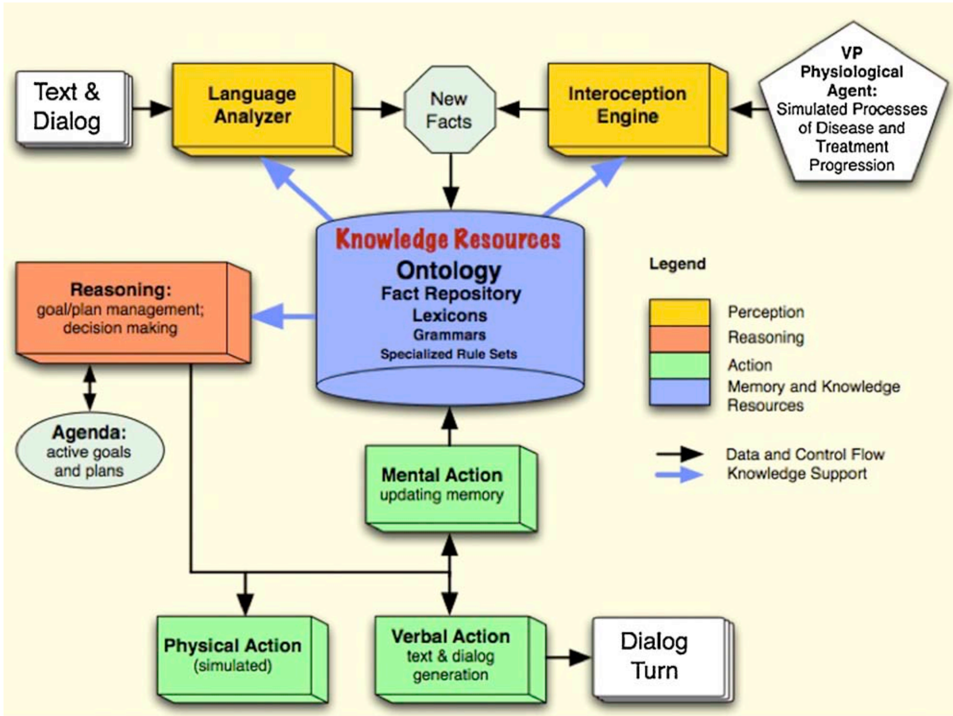

Fig. 1. The architecture of "double agents" in OntoAgent. The physiological agent is represented by the upper right pentagon. The rest of the figure – including the interoception engine, which perceives physiological stimuli generated by the physiological agent – is the realm of the cognitive agent.

either unknown or unnecessary for the needs of realistic simulation [1,2].

- Modeling interoception, which is the cognitive perception of physical signs and symptoms by intelligent agents [3].
- Decision theory, including goal and plan management and hybrid (rule-based, analogy-based and probabilistic) reasoning [3].
- Deep language processing that includes semantic and pragmatic analysis and language generation [4,5].
- Dialog modeling [6].
- Memory modeling and management [7–9].
- Agent individualization according to character traits (courage, suggestibility, boredom threshold etc.), personal preferences (likes coffee, is afraid of surgery, etc.), differing knowledge of the world, etc. [10].
- Agent learning of facts, concepts and language by being told, by reading and by experience [11].
- Use of a shared metalanguage of description and knowledge bases for agent reasoning, learning and language processing [12,13].
- Semi-automatic knowledge elicitation for knowledge-based systems [14].
- Integrating complex multipurposes knowledge systems.
- Ergonomic issues for developers and subject matter experts, such as the development and maintenance of a uniform knowledge representation substrate for core knowledge bases and processors and the development of a variety of efficiency-enhancing toolkits ([4] and http://www.trulysmartagents.org/dekade.php).
- Validation of our work by demonstrating feasible, proof-of-concept applications [1,2,10,11,15].

Here, we provide brief background about select aspects of OntoAgent that should suffice to support understanding of the new work being reported.

Intelligent agents in OntoAgent are "double agents", in that they have a cognitive side and, optionally, a physiological side. The architecture of agents is shown in Fig. 1.

To model the physiological agent – which "lives" a simulated life over time in applications – we encode knowledge about biophysical functions that have clinical relevance in the maintenance of health, the production of disease, and the bidirectional transitions between these two states. When biomechanisms are known and are deemed clinically important, they are modeled using causal chains. Gaps in our knowledge of biomechanisms are bridged with non-biomechanistic knowledge from the literature, practical clinical knowledge and, in some cases, probabilistic methods. The depth and granularity of the models are determined by the demands of automatic function and realism in our current and anticipated applications. As a rule of thumb, a feature value or process is included in the model if it can either be measured through tests, be affected by medication/interventions, or cause a change in some other clinically relevant feature value or process. Of central importance is the fact that our models can be easily modified, reflecting new medical findings or the assessment that a clinical "bridge" (i.e., a non-mechanistic observation of how a disease process manifests based on population-level clinical observation) needs to be replaced by a biomechanism in order to support the functionality required by an application.

The cognitive agent engages in the well-known triad of functionalities: perception, reasoning and action. It is ideologically close to, though methodologically not identical with, the belief-desire-intention (BDI) model of agency [16]. Unlike the classical BDI implementations [17], our approach centrally involves language comprehension and production as well as physiological simulation.

As shown in Fig. 1, agents undergo two types of *perception*: interoception and language understanding. The results of both modes of perception are interpreted and stored in the agent's knowledge resources: ontology, fact repository (memory of assertions), and language resources, such as lexicon and grammar (the latter can be learned, e.g., when the agent encounters new words through interaction with another agent). Each agent has its own knowledge resources, which can differ from those of other agents to represent different levels of education, different personal experience,