Cognitive and Metacognitive Educational Systems: Papers from the AAAI Fall Symposium (FS-10-01)

Aspects of Metacognitive Self-Awareness in Maryland Virtual Patient

Sergei Nirenburg, Marjorie McShane and Stephen Beale

Institute for Language and Information Technologies
University of Maryland Baltimore County
1000 Hilltop Circle, Baltimore, MD 21250
{sergei,marge,sbeale}@umbc.edu

Abstract

This paper describes Maryland Virtual Patient (MVP), a simulation and tutoring environment developed to support training cognitive decision making in clinical medicine. MVP is implemented as a society of agents, with one role – that of the trainee – played by a human and other roles played by artificial intelligent agents. In order to make the trainee's experience as similar as possible to the traditional medical training environment, MVP is implemented as a collection of knowledge-based models of simulated human-like perception, reasoning and action processes. MVP operation involves metacognition: for example, the MVP virtual patient is aware of the physiological state of its body, of its physiological and character traits as well as of lacunae in its knowledge about the world and about language. This self-awareness influences the virtual patient's reasoning and actions. In this paper we illustrate the role of metacognitive self-awareness in the overall operation of MVP.

Introduction

The Maryland Virtual Patient (MVP) system (Figure 1) is a simulation and tutoring environment developed to support training cognitive decision making in clinical medicine. MVP is implemented as a multi-agent society with one human agent (the trainee) interacting with several artificial agents – the VP, other medical personnel (lab technicians, specialists, etc.) and a tutor (McShane 2007a, b).

Virtual patients (VPs) in MVP are "double agents" combining a physiological agent that simulates the progression of a disease with a cognitive agent (Figure 2) capable of perceiving symptoms, understanding natural language inputs, reasoning (notably, making decisions about its actions) and carrying out verbal actions. The language capabilities allow the VPs not only to engage in dialog but also to learn by being told. The VPs can learn both general facts (for example, the properties of a disease) and specific facts (for example, that their test results are negative).

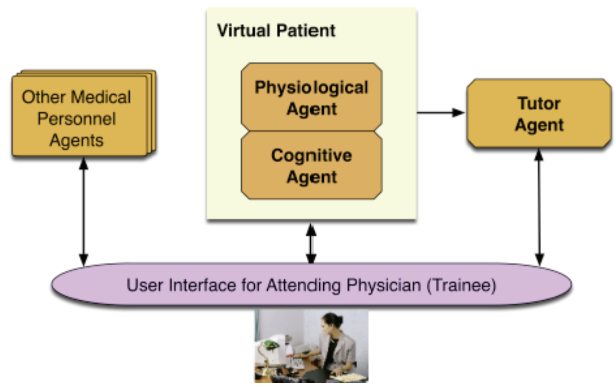

Figure 1. The agents in the MVP environment.

The cognitive side of the VP involves many aspects of metacognition, including: the VP's understanding of its own physical state; its understanding and manipulating of decision spaces; its understanding of its own goals and its opinions about how best to achieve them; and its utilization of natural language to convey its thoughts and understand the thoughts of others. In the next section, we discuss each of these aspects of metacognition in turn, with examples from the MVP system. We then provide a screen-shot supplemented walk through system capabilities, previewing what we will demonstrate at the conference.

Five Aspects of Metacognition

1. Interoception: Perceiving the body's signals

Interoception is the perception of physiological phenomena. It is a VP feature that has both physiological and cognitive aspects. The source of interoception is physiological phenomena, like symptoms of a disease, hunger and sleepiness. The VP experiences current symptoms of its disease and has memories of previous symptoms, including their severity, so that useful comparisons can be made:

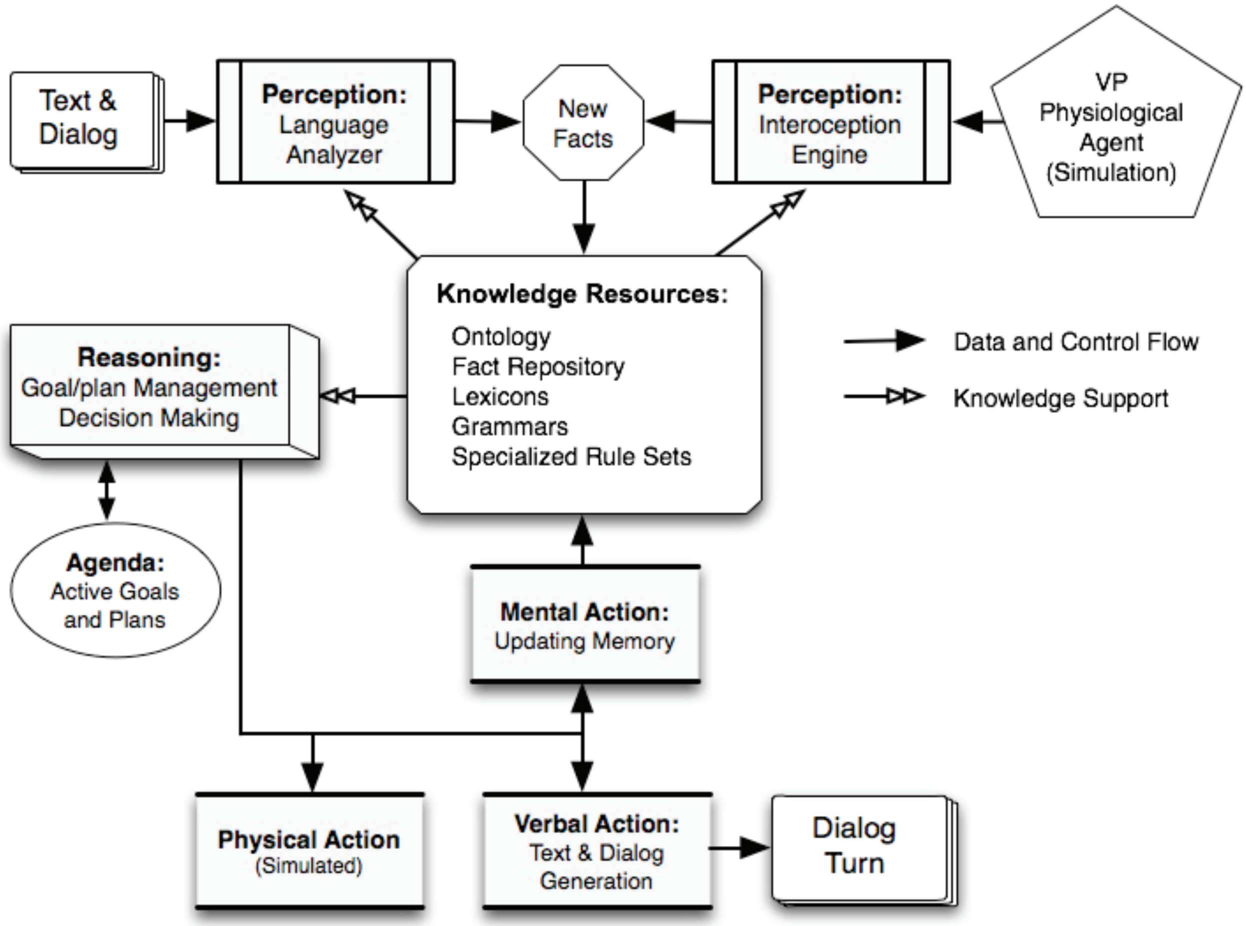

Figure 2. The architecture of MVP.

e.g., “Symptom X has gotten much worse over the past month, I had better go see my doctor sooner than our next scheduled appointment.”

Memories are stored using an ontologically grounded metalanguage that is identical to the one used to represent the meaning of language input (cf. below). For example, the property called “health-attribute” generalizes over the VP’s symptoms so that it can assess its overall health, which is an input to certain kinds of decision-making. Figure 3 shows different values of health-attribute at different times during a simulation run of a VP: the lower the value, the worse the perceived health and the more likely the patient will go to the doctor.[1] Of course, when memories about interoception are stored, there need be no translation into and from a natural language: the entire process occurs at the level of the metalanguage.

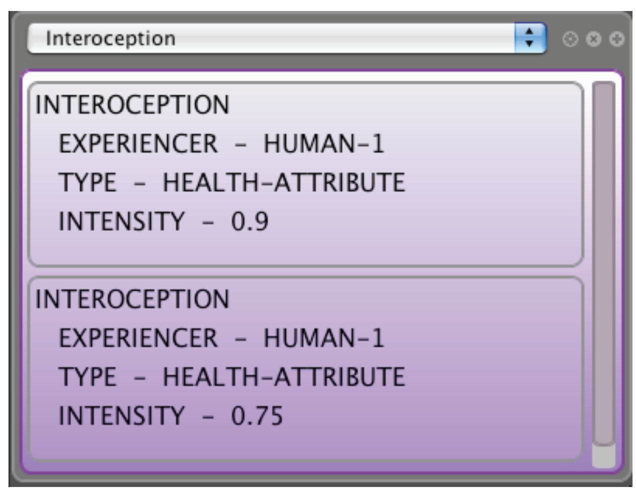

Figure 3. Results of interoception at different times of the VP’s disease progression.

The experiencing of symptoms is individualized for each VP instance through the use of character traits (e.g., trust in the doctor) and physiological features (e.g., how well the patient tolerates treatments). When a given VP instance is created, values for these features are selected and affect the VP’s reactions in the face of its disease(s). Values for the physiological aspects of the disease(s) and the VP’s re-

[1] The panel shown in Figure 3 is one of the “under-the-hood” views of MVP available in the system to show select traces of system functionality.

sponse to interventions, should they be applied at various times, are also selected for each individual VP.

2. Understanding and Manipulating a Decision Space

An important feature of modern medical training is involving the patient in his or her own health management. To facilitate this in a simulated environment with a virtual patient we chose to endow the VP with the capability of dynamic (not preprogrammed) decision making that approximates human decision making (Nirenburg et al. 2008). Whenever a decision needs to be made, the VP first determines whether it has sufficient information to make the decision. This initial assessment is based on a combination of what it actually knows, what it believes to be necessary for making a good decision, and its personality traits. If the VP concludes that it lacks some decision-making knowledge, it posits the goal of obtaining this knowledge. This behavior is clearly metacognitive and involves a learning process initiated by the VP.

For example, say the trainee suggests that the patient have the endoscopic procedure pneumatic dilation, which the VP has never heard of. If the VP were very trusting it might agree to the procedure then ask what it was, or it might agree to it and wait to see if the doctor provided any information about – if not, it might just go home and look it up on the Web. If the VP were not very trusting, or if it were trusting but just very curious, it might ask questions before agreeing. The VP knows what questions to ask – i.e., what features to ask about – based on (a) its ontological knowledge about objects and events in the world (the properties each one *can* have) and (b) the VP's own personality traits, which makes it interested in some subset of those properties. For example, all of our VPs know that medical procedures can be painful, they can carry risks, etc.; so if a VP is particularly worried about pain, but is very courageous when it comes to risks, then it will ask about the pain but be unconcerned about the risk level and leave that feature value unspecified in its ontology even at the point of decision-making.

3. The VP's awareness of its goals, and its opinions about how to best achieve them

In some aspects of human behavior – e.g., during explicit decision-making: choose X or Y – people tend to be aware of their goals and preferences and how they affect their behavior. However, there exist other kinds of contexts in which a person's goals affect his behavior, whether or not the person is aware of it. One example is a patient's choice of how to convey its subjective experiencing of illness (Searight and Campbell 1992) to the doctor. One patient experiencing moderate pain might convey it to the doctor as a severe symptom (e.g., because he or she craves attention from the medical community) whereas another might convey it as a mild symptom (e.g., because it is considered not honorable to complain about pain or because he/she does not have the money to pay for treatment anyway). Practicing clinicians must be aware of this aspect of metacognition – i.e., interpretation of symptoms – in order to effectively diagnose and treat patients. As such, artificial VPs must show a great variety of behaviors in this area, displaying both the normal case, in which symptoms more or less correlate with the known aspects of the disease state, and the atypical case, when symptoms either do not correlate or patients are misrepresenting them, consciously or not. VPs in MVP have character traits that permit them to under- and over-represent their symptoms to provide for this kind of variety and the diagnostic challenges it poses.

4. Natural Language as the Medium of Conveying One's Thinking

We view natural language use as metacognition because it is the medium by which humans can convey their own thinking and gain insights into the thinking of others. We have chosen to enhance verisimilitude in language interaction through increased modeling sophistication of language use, but this carries a steep price tag. As an illustration, compare the amount of work carried out by the MVP system and a representative state-of-the-art VP system (Chesher 2004) to respond to a dialog turn of the user. The latter system relies on matching keywords from the input with a list of inputs for which it has prepared answers (dialog turns to be generated by the system). By contrast, in order to generate a comparable dialog turn, the VP in the MVP environment has to perform a large number of non-trivial perception (text understanding), reasoning (goal- and plan-directed decision making) and action-related (verbal action) operations. For example, to answer the "simple" query *How are you feeling?*, the VP must have the -capability to: a) extract the meaning of the user's dialog turn, including its illocutionary force, i.e., the speech act meaning; b) add the resulting text meaning representation to the short-term memory component of its fact repository; c) generate an instance of a "Be-a-Cooperative-Conversationalist" goal; d) prioritize goal instances on its agenda; e) select a goal instance for processing; f) [once the above goal instance is chosen] select a plan to pursue to attain this goal (there is currently just one plan for this type of goal: "carry out a relevant verbal action"); g) specify the content of the verbal action to be produced; and h) generate an English sentence that realizes the above content, which is a report about its health attribute. In the current version, lexical selection for the answer to this question is based on the value of health-attribute and syntactic structure selection is analogy-based, driven by random selection from an inventory of sentence patterns. In sum, language processing by agents in MVP aims to extract deep semantics from an interlocutor's utterances, translate the meaning into the same metalanguage used for all knowledge representation

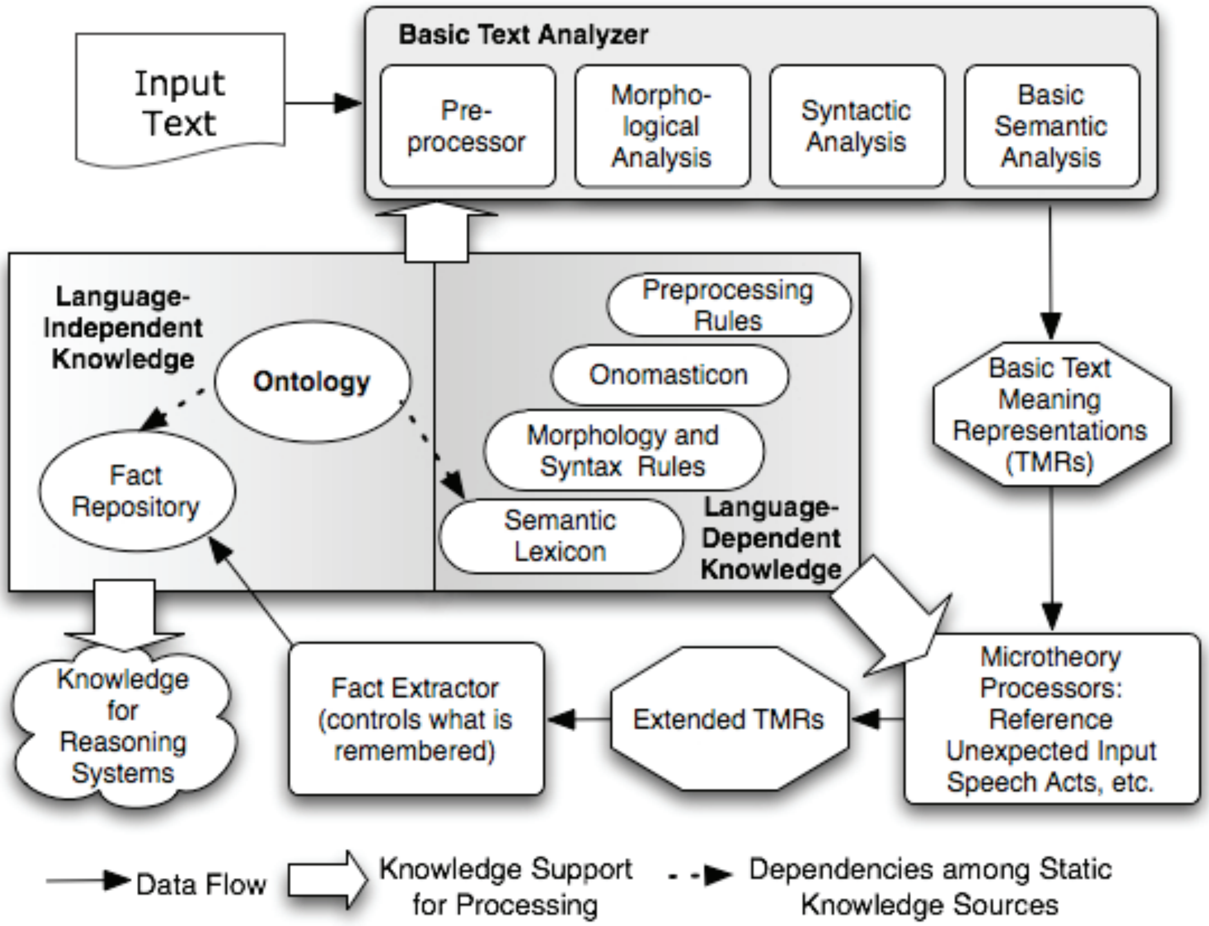

Figure 4. The OntoSem semantic analyzer.

in the system, and thereby support truly high-level reasoning.

MVP System Demonstration

MVP uses a library of VPs created by instructors using a specially designed interface (Jarrell et al. 2008). An MVP session starts with the teacher loading a VP from the library of VPs. The user (the trainee) then starts the VP (the physiological agent) simulation. The progression of the simulated disease leads to the perception of symptoms by the cognitive agent of the VP, which eventually causes the VP to seek medical help. The simulated actions include making an appointment and showing up at the doctor's office. Once the VP has come in for a visit, the trainee starts a dialog, with the goal of diagnosing and treating the VP.

The trainee must ask diagnostic questions, establish a diagnosis, suggest diagnostic and, at a later stage, treatment procedures, answer the VP's questions about the diagnosis and about the suggested tests and treatments, and discuss alternative tests and treatments. The VP can initiate unscheduled visits if its state of health rapidly declines.

The MVP interface that is used by the trainee is illustrated in Figure 5. The system shows the dialog as well as notes that it automatically makes in the online chart of the VP (accessible by the tab to the left).

Figure 5. The MVP trainee interface.

While it cannot be guaranteed that the VP agent will always understand the meaning of a user input, the system is capable of generating correct text meaning representations for a wide variety of inputs. The OntoSem analyzer relies on a lexicon of over 30,000 word senses (which are not constrained to the application domain) that are interpreted in terms of an ontology containing about 9,000 concepts characterized on average by 16 properties each.

Figures 6-9 show four individual panes that, like Figure 3, are part of MVP's "behind the scenes" interface that is used for explaining and demonstrating the processing carried out by the VP as well as for testing and debugging the system. Figure 6 shows a partial view of the anatomic elements of the VP (the esophagus, the lower esophageal sphincter, etc.) and values of their properties at a particular moment in the progression of VP's disease (achalasia, a disease of the esophagus). When certain property values reach predetermined thresholds, an instance of interoceptive perception is triggered, as a result of which the VP obtains meaning representations illustrated in Figure 3.

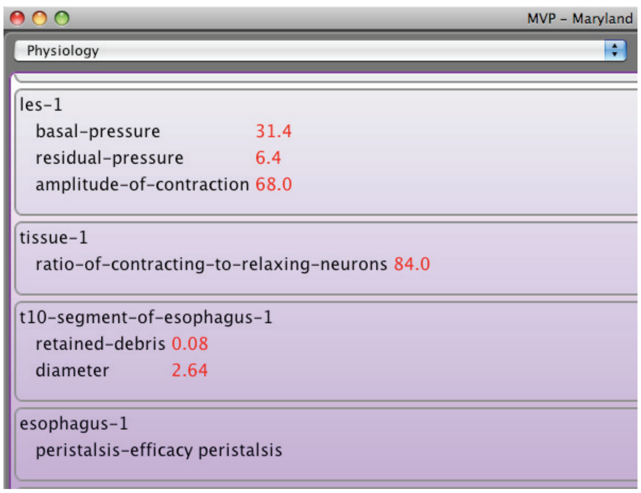

Figure 6. A window into the VP's physiology.

As can be seen in the figure, health-attribute, the main overall measure of the VP's state of health, deteriorates over time. Note that the format of these representations is identical with that of the text meaning representations created when the VP receives language input. An example of a text meaning representation is shown in Figure 7 – it represents the input "*How's the swallowing?*" (or any of the paraphrases of that locution; see McShane et al. 2008 for more on paraphrase).

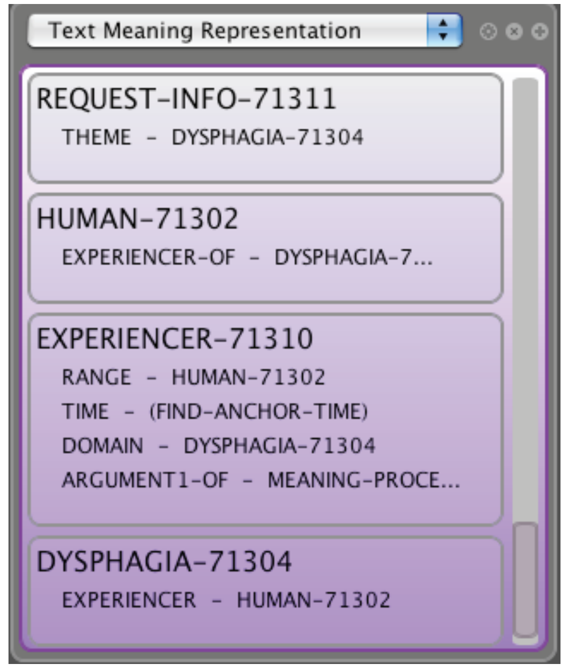

Figure 7. A Sample TMR.

Figure 8 is a window into the decision making process of the VP. For ease of observation, the decisions are shown in the panel as English renderings rather than the metalanguage structures that are actually manipulated by the VP. The first two "thoughts" are from before the VP comes in for an office visit. The rest of the thoughts reflect the VP's decisions for the following portion of the dialog:

Trainee: I suggest that you have an EGD, which is a diagnostic procedure.
VP: How risky is it?
Trainee: It's not risky.
VP: Are there any side effects?

The final "under the hood" panel we will show (Figure 9) traces the results of the VP's learning. In the above dialog, the VP heard the term EGD for the first time. From analyzing the text, it was able to create a new (for it) ontological concept EGD and lexicon entry EGD-N1 (that is, first nominal sense of EGD).

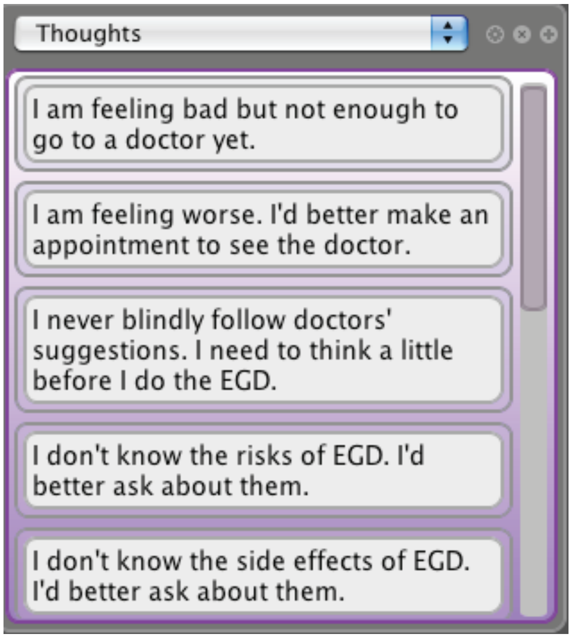

Figure 8. The VP's "Thoughts."

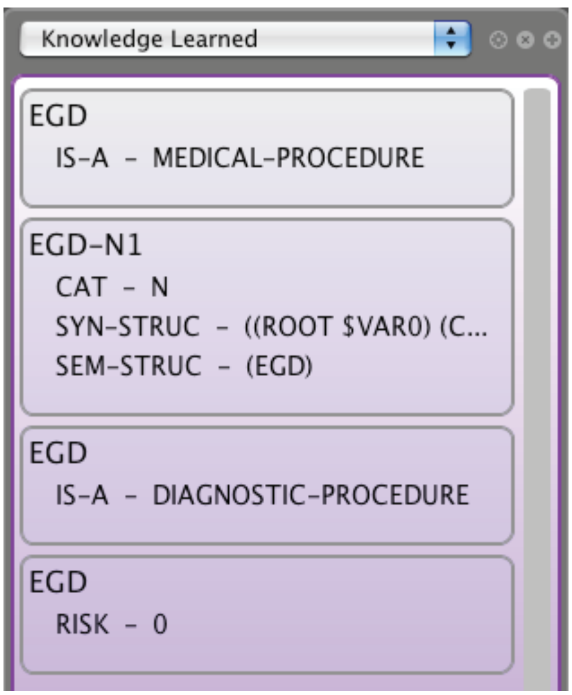

Figure 9. The VP learns.

At the conference we will show live demos of creating, diagnosing and treating virtual patients. More information about the MVP project (including a video of a demo and bibliography) can be seen at www.trulysmartagents.org.

References

Chesher, D. 2004. Exploring the use of a web-based virtual patient to support learning through reflection. Ph.D. diss., University of Sydney.

Jarrell, B., Nirenburg, S., McShane, M., Fantry, G. and Beale, S. 2008. Revealing the conceptual substrate of biomedical cognitive models to the wider community. *Medicine Meets Virtual Reality 16*. Newport Beach, Calif.

McShane, M., Fantry, G., Beale, S., Nirenburg, S. and Jarrell, B. 2007a. Disease interactions in cognitive simulations for medical training. In Proceedings of the MODSIM World Conference, Medical Track. Virginia Beach, VA.

McShane, M., Nirenburg, S., Beale, S., Jarrell, B., and Fantry, G. 2007b. Knowledge-based modeling and simulation of diseases with highly differentiated clinical manifestations. In Proceedings of the 11th Conference on Artificial Intelligence in Medicine. Amsterdam, Netherlands.

McShane, M., Nirenburg, S. and Beale, S. 2008. Two kinds of paraphrase in modeling embodied cognitive agents. In Proceedings of the Workshop on Biologically Inspired Cognitive Architectures, AAAI 2008 Fall Symposium. Arlington, VA.

Nirenburg, S., McShane, M. and Beale, S. 2008. A simulated physiological/cognitive "double agent". In Proceedings of the Workshop on Naturally Inspired Cognitive Architectures at AAAI 2008 Fall Symposium. Arlington, VA.

Nirenburg, S. and Raskin, V. 2004. Ontological Semantics. Cambridge, Mass.: The MIT Press.

Searight, H.R., and Campbell, D.C. 1992. Ethnography and family medicine: Issues and overview. The Family Practice Research Journal 12(4): 369-382.

Copyright @ 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.