Integrating Cognitive Simulation into the Maryland Virtual Patient

Sergei NIRENBURG[a], Marjorie MCSHANE[a], Stephen BEALE[a], Bruce JARRELL[b] and George FANTRY[b]
[a]*University of Maryland Baltimore County, Baltimore, Maryland*
[b]*University of Maryland School of Medicine, Baltimore, Maryland*

Abstract. This paper briefly describes four cognitively-related aspects of modeling a virtual patient: interoception, decision-making, natural language processing and learning. These phenomena are treated within the Maryland Virtual Patient simulation and training environment.

Keywords. simulation, cognitive simulation, physiological simulation, medical training

1. Overview

The Maryland Virtual Patient[1] (MVP) project is developing an agent-oriented environment for automating certain facets of medical education that includes a network of human and software agents. The human agents include the user (typically, a trainee) discharging the duties of an attending physician and, optionally, a human mentor. The human-like software agents include the virtual patient, lab technicians, specialist consultants and a mentoring agent. The system also includes an array of non-humanlike software agents.

At the core of this network is the virtual patient – a knowledge-based model and simulation of a person suffering from one or more diseases [1-4]. The virtual patient is a "double agent" in that it models and simulates both the physiological and the cognitive functionality of a human. Physiologically, it undergoes both normal and pathological processes in response to internal and external stimuli. Cognitively, it experiences symptoms, has lifestyle preferences (a model of character traits), has memory (many of whose details fade with time), and communicates with the human user about its personal history, symptoms and preferences for treatment.

Users can interview a virtual patient; order lab tests; receive the results of lab tests from technician agents; receive interpretations of lab tests from consulting physician agents; posit hypotheses, clinical diagnoses and definitive diagnoses; prescribe treatments; follow-up after those treatments to judge their efficacy; follow a patient's condition over an extended period of time, with the trainee having control over the speed of simulation (i.e., the clock); and, if desired, receive mentoring from the automatic mentor.

---

[1] Patent pending.

The virtual patient (VP) simulation is grounded in an ontologically-defined model of human anatomy and physiology. Instances of virtual patients with particular diseases and particular physiological traits are generated from ontological knowledge about human physiology and anatomy by grafting a disease process onto a generic instance of a human. Disease processes themselves are described as complex events in the underlying ontology.

We have reported progress on this project at past MMVR meetings. Specifically, we described our approach to physiological simulation [3] and the way we can succinctly describe the main features of each disease model to developers, other experts and the wider community [4]. We continue this year by discussing newly integrated cognitive capabilities of the virtual patient. The cognitive side of the VP currently models several aspects of cognitive processing:

- **interoception** – the perception of physiological phenomena, such as symptoms, and the interpretation and remembering of such phenomena

- **decision making**
  - deciding when to go see a physician, both initially and during treatment
  - deciding whether to seek help in making decisions related to treatment by asking the user knowledge-seeking questions about a recommended test or intervention
  - deciding whether to agree to a recommended test or intervention

- **natural language processing**
  - language perception and understanding, including both the direct meaning of physician-user communication in natural language and its intent
  - deciding on what specifically to communicate to the user
  - actually generating natural language utterances

- **learning** – receiving new knowledge about the world and the words and phrases used to describe it and adding them to the ontology and lexicon, respectively

We will briefly touch upon each of the above points with the goal of describing *what* the virtual patient can do rather than *how* it can do it. A sufficient discussion of the latter would require far more space.

2. Interoception

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. Here we focus on physiology symptoms because our disease models to date have not required the tracking other kinds of interoception.

The VP experiences current symptoms of its disease and has memories of previous symptoms, including their severity, so that useful comparisons can be made: 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). 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.

The experiencing of symptoms is individualized for each VP instance through the use of character traits and physiological features. Our current inventory of character traits includes trust (trust in the doctor’s advice), suggestibility (how readily the VP agrees to the doctor’s recommendations) and courage (how willing the VP is to undergo tests or procedures even if they are risky or have significant side effects). Our current inventory of physiological traits includes physiological-resistance (e.g., how well the patient tolerates treatments), pain threshold (how much pain the VP can stand) and the ability to tolerate symptoms (how intense or frequent symptoms have to be before the VP consults a doctor). This inventory clearly must be expanded in the future. When a given VP is created, values for these features are selected and affect the VP’s reactions in the face of its disease(s). Of course, values for the physiological aspects of the disease(s) and the VP’s response to interventions, should they be applied at various times, are also selected for each individual VP.

3. Decision Making

The decision-making behavior of specific instances of virtual patients is parameterized using a model of personality traits and physical and mental states. It is informed by (a) the content of the VP’s short-term memory, which is modeled as knowledge invoked specifically for making the decision at hand, and (b) the content of the VP’s long-term memory, which is the VP’s recollection of its past states of health, past communications and decisions, and general world knowledge.

A VP’s decision-making, as described above, is affected by the severity and duration of its symptoms; its knowledge of tests and procedures; the character traits trust, suggestibility and courage; and the physiological traits physiological-resistance, pain-threshold and the ability to tolerate symptoms.

VP reasoning is carried out through modeling the VP’s goals and plans, thus broadly conforming to the belief-desire-intention (BDI) approach to developing intelligent agents [5] (see also [6] for related discussion).

When the VP starts to experience symptoms it can either do nothing, go the doctor, go to the emergency room or self-treat. As concerns seeing a doctor for the first time, the VP compares its symptom severity with its ability to tolerate symptoms – as well as some character traits not yet incorporated, such as its desire to be fussed over by a doctor vs. its dislike of seeing doctors. Later in its treatment the VP also considers the date of its next scheduled appointment, whether or not its symptoms have spiked, etc.

When the doctor recommends a test or procedure, the VP must compare its knowledge of the test/procedure with its character traits, like courage and suggestibility. For example, if it knows nothing about the test/procedure and has little trust in the doctor, it will ask questions about the properties that interest it, like the pain level and side effects; by contrast, if it has complete trust in the doctor and a high value for suggestibility, it will ask no knowledge-seeking questions and, instead, agree to anything the doctor suggests.

When the patient has received all the knowledge of tests/procedures it feels it needs, it will decide whether or not to agree to the test/procedure. It can also suggest other options that it happens to know about, and the doctor can accept or reject such suggestions.

4. Language Processing

Our approach to treating language communication is unlike most other approaches in that all language-oriented reasoning is carried out on the basis of formal interpretations of the meaning of linguistic expressions. Our automatically generated, semantically-oriented text meaning representations (TMRs) are written using the same ontological knowledge substrate and the same ontologically grounded metalanguage as are used to represent physiological processes, interoception and agent goals and plans. In short, all knowledge and reasoning in our environment employ the same metalanguage, so whether a VP experiences new symptoms (through interoception) or learns information about its disease from the user (through language processing), the new information will be stored the same way in the VP's memory.

There are several advantages to orienting an agent's language processing around TMRs rather than text strings. First, TMRs are unambiguous, since linguistic ambiguity is resolved as the TMRs are being generated. Second, TMRs reduce to a single representation many types of linguistic paraphrase, be it lexical (*esophagus ~ food pipe)*, syntactic (*I will administer it to you ~ It will be administered to you by me*) or semantic (*Does the food get stuck when you swallow? ~ Do you have difficulty swallowing?*) [7]. Third, TMRs facilitate the detection of which aspects of meaning are central and which are of secondary importance. As regards paraphrase processing, in addition to having to resolve linguistic paraphrase, the VP must be able to resolve two other kinds of paraphrase: a) the reformulation of the representation of physiological events (e.g., symptoms) in "lay" ontological terms that can be understood and remembered by its cognitive agent and b) the representation of the meaning of verbal messages in terms compatible with how related content is stored in the cognitive agent's memory [8]. Remember, the VP is typically not a medical professional, meaning that it must have a different ontology and a different lexicon than a physician would.

5. Learning

We just noted in passing that the VP's ontology and lexicon do not match those of a physician. Indeed, the physician's ontology will contain a vast subtree of medical information including objects, events and the properties that link them as well as script-based knowledge [9], which permits the physician to understand the progression of a disease, how to treat it under various circumstances, etc. The physician will have a correspondingly large technical and non-technical vocabulary (lexicon) that is linked to the respective ontological concepts and is used to analyze and generate language in the medical domain. The patient's knowledge base, by contrast, will typically include an impoverished medical subtree in the ontology and a relatively small number of medical terms in the lexicon – unless, of course, the VP happens to be a physician or even a specialist in the given domain, which raises a different set of complications for the user.

In conducting an interview with a VP, the user must be able to express himself in different ways, using paraphrases selected according to the degree of medical knowledge the VP possesses. However, during appointments the physician will naturally teach the VP about various aspects of its condition: its name, the names of related drugs and procedures, the properties of drugs and procedures that the VP asks about or the user chooses to provide, the medical terms for words that formerly had to be paraphrased for the VP, and so on. For example, (a) when the user tells the VP the name of its disease, that disease is added to the VP's ontological subtree of diseases and a new lexicon entry is created that maps to this ontological concept; (b) when the VP learns information about a test or procedure it remembers it and no longer asks questions about it – unless, of course, the VP forgets, in which case the user will need offer a reminder. (See related discussions in [10-12].)

6. Discussion

This paper has briefly described four cognitive aspects of VP modeling: interception, decision-making, natural language processing and learning. Although each of these presents significant challenges, the work has been facilitated by using the same ontological substrate for all aspects of modeling.

MVP is a classical AI (artificial intelligence) system in that it strives to model human perception, reasoning and action capabilities and does so on the basis of encoded knowledge. It differs from much of classical AI practice in that it includes people as components in its architecture. If MVP were an expert system in the classical sense, *the system* would have been tasked to diagnose and treat patients rather than the other way around. Indeed, many systems in the medical domain, from Mycin [13] on up, had this as their main goal. Another difference from classical AI is the centrality of the descriptive component of the system: the VP's world is certainly not toy (i.e., covering an extremely small domain). Our emphasis is on acquiring knowledge that is sufficiently deep to support the complex reasoning, simulation and language processing required by the application. This is in contrast to many recent and current approaches (notably, in natural language processing) that stress broad coverage of data in contrast to the utilization of a depth of acquired knowledge.

We believe that the statements often heard nowadays about the demise of AI are ill conceived. The AI enterprise did not fail. In fact, it has not yet been brought to the test. This is because the enterprise is much more complex than it was perceived to be even by many AI practitioners themselves. Despite the recent emphasis in the field on statistics-oriented methods, they should not be viewed as having superceded classical AI. In fact, these methods have contributed to the core task of knowledge acquisition that is a prerequisite to the success of the program of AI. Progress in learning, knowledge visualization and other ergonomic factors, the ease of access to vast collections of data on the Web and other developments make the original AI goals incrementally more attainable. Our work on the MVP corroborates this state of affairs.

We believe that the development of a comprehensive MVP is feasible both scientifically and logistically.

References

[1] M. McShane, S. Nirenburg, S. Beale, B. Jarrell, G. Fantry, Knowledge-based modeling and simulation of diseases with highly differentiated clinical manifestations, *Proceedings of the 11th Conference on Artificial Intelligence in Medicine (AIME 07)*, Amsterdam, The Netherlands, July 7-11, 2007.

[2] M. McShane, G. Fantry, S. Beale, S. Nirenburg, B. Jarrell, Disease interaction in cognitive simulations for medical training, *Proceedings of the MODSIM World Conference, Medical Track*, Virginia Beach, Sept. 11-13, 2007.

[3] B. Jarrell, S. Nirenburg, M. McShane, G. Fantry, S. Beale, D. Mallott, J. Razcek, An interactive, cognitive simulation of gastroesophageal reflux disease, *Stud Health Technol Inform* (2007) 194 –199, IOS Press.

[4] B. Jarrell, S. Nirenburg, M. McShane, G. Fantry, S. Beale, Revealing the conceptual substrate of biomedical cognitive models to the wider community, *Stud Health Technol Inform* (2008), 281 – 286, IOS Press.

[5] A.S. Rao, M.P. Georgeff, Modeling rational agents within a BDI-architecture, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann Publishers, Inc., San Mateo, CA, USA, 1991, 473 – 484.

[6] S. Nirenburg, M. McShane, S. Beale, A simulated physiological/cognitive "double agent", *Proceedings of Naturally Inspired Cognitive Architectures, AAAI 2008 Fall Symposium*, Washington, D.C., Nov. 7-9, 2008 (forthcoming).

[7] M. McShane, S. Nirenburg, S. Beale, Resolving paraphrases to support modeling language perception in an intelligent agent, *Proceedings of the Symposium on Semantics in Systems for Text Processing* (*STEP 2008*), Venice, Italy, 2008.

[8] M. McShane, S. Nirenburg, S. Beale, Two kinds of paraphrase in modeling embodied cognitive agents, *Biologically Inspired Cognitive Architectures, AAAI 2008 Fall Symposium*, Washington, D.C., Nov. 7-9, 2008 (forthcoming).

[9] R. Schank, R. Abelson, *Scripts, Plans, Goals and Understanding*, Hillsdale, NJ, Erlbaum, 1977.

[10] S. Nirenburg, T. Oates, Learning by reading by learning to read, *Proceedings of the International Conference on Semantic Computing* (*ICSC-07*), Irvine, CA, 2007.

[11] P.Z. Yeh, B. Porter, K. Barker, Using transformations to improve semantic matching, *Proceedings of the 2nd International Conference on Knowledge Capture*, Sanibel Island, FL, USA, 2003, 180–189.

[12] S. Nirenburg, M. McShane, S. Beale, B. Jarrell, Adaptivity in a multi-agent clinical simulation system, *Proceedings of AKRR'08, International and Interdisciplinary Conference on Adaptive Knowledge Representation and Reasoning*, Porvoo, Finland, Sept. 17 – 19, 2008.

[13] B. Buchanan, E. Shortliffe, *Rule-Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley, 1984.