Cognitive Modeling for Clinical Medicine

Sergei Nirenburg and Marjorie McShane

Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
{marge,sergei}@umbc.edu

Abstract

This paper describes some functionalities and features of the Maryland Virtual Patient (MVP) environment. MVP models the process of disease progression, diagnosis and treatment in virtual patients who are endowed with a "body," a simulation of their physiological and pathological processes, and a "mind," a set of capabilities of perception, reasoning and action that allow the virtual patient to exhibit independent behavior, participate in a natural language dialog, remember events, hold beliefs about other agents and about specific object and event instances, make decisions and learn.

Background

Clinical medicine is a collaboration between medical practitioners (MPs) and patients: after all, many of the best diagnostic and treatment plans will have less of a chance of succeeding if the patient does not comply and contribute. Over time MPs develop a wide range of skills related to understanding what a patient might leave unsaid, predicting what aspects of life and health outside of the chief complaint should be discussed, anticipating when a patient might not tell the truth in order to optimize something outside of the treatment of his chief ailment, educating the patient to assuage his fears and/or convince him to agree to a course of treatment, preparing backup plans if the patient refuses some aspect of best treatment practices, and so on. Ideally, this knowledge – which currently comes only with years of experience caring for real patients – could be acquired more quickly. This would be possible if a MP could interact with and manage the care of a large population of physiologically distinct virtual patients who also differed widely in personality traits, culturally-, socially- and religiously-oriented beliefs, and negotiation skills – all of which affect their decision-making about medical care. The above desiderata are at the moment infeasible in medical education because of the limitations inherent in current medical school curricula:

- Future physicians do not have the opportunity to adequately practice their cognitive analysis and problem solving skills, especially with regard to treating non-emergency patients over long periods of time.
- During their apprenticeships, future physicians typically do not see a) patients with all the diseases that must be studied or b) a sufficient number of diverse cases of a particular disease.
- There is not enough time for teachers to spend with individual students and it is economically infeasible to hire more teachers.

One way of alleviating this state of affairs is to teach future physicians by having them diagnose and treat virtual patients (VPs). Several kinds of VPs have been developed and are relatively widely used in medical training, among them technical task trainers for developing motor skills and fully preprogrammed narrative-oriented interactive scenarios for developing cognitive skills.

The Maryland Virtual Patient (MVP)[1] project represents a departure from state-of-the-art medical training systems because it models VPs as cognitive agents with perception, reasoning and action capabilities, who are also endowed with models of character traits and personal preferences, memory and learning. These VPs thus require extensive knowledge resources that cover knowledge about the world, about other agents in that world and about language. (A peculiar characteristic of the MVP environment is the obvious need to treat the agent's own body as part of the physical world in which a VP is immersed.) Improving all of the above capabilities and enhancing the knowledge substrate of the agents is the long-term goal of the project. It comes as little surprise that this goal clearly dovetails with the core goals of cognitive science and artificial intelligence. As such, it comes with equally little surprise that, while our results to date are very encouraging, much more work than has been accomplished so far will have to be done to attain the desired breadth of coverage and perception of adequate verisimilitude.

[1] Patent pending.

To make VPs more sophisticated, The MVP project has pursued four interleaved directions of work: a) simulating normal and pathological physiological processes; b) modeling human cognitive, including conversational, capabilities; c) formalizing the clinical experience of master physicians and using the knowledge of these "best clinical practices" to support automatic tutoring; and d) integrating the VP in a computational environment modeling a society of intelligent agents.

We believe that developing simulated cognitive capabilities of artificial intelligent agents is essential for many high-end applications, even taking into account the high cost involved. We also believe that attaining this level of sophistication on a realistic scale is feasible. In our opinion, the currently prevalent empirical, statistics-oriented methods should be viewed not as a substitute for classical AI but as useful tools for eliminating the knowledge acquisition bottleneck that has so far slowed AI down. 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 experience in the MVP project so far seems to support this opinion. In the rest of this paper, we will briefly present some of the VP functionalities.

The Physiological Agent

The subject domain of the MVP project is clinical medicine. The core tasks in building the physiological agent are a) developing clinical theories of disease progression and b) implementing the theoretical parameters and their fluctuations in a knowledge-based simulation environment. We call our theories clinical to underscore the fact that many of the physical, chemical and physiological processes and their causal connections are either not understood or are not necessary to support realistic simulation. Whenever such processes are encountered in describing disease progression, we use "bridges," elements of clinical expertise that distill expert knowledge of medicine. The resulting model, thus, constitutes a medically correct approximation of a truly causal model of disease progression.

The availability of a simulated physiological model allows us to:

- forgo specific authoring of a VP's state of health at a particular time, including effects and side effects of interventions and treatments
- permit a VP to respond realistically to "unexpected" interventions, that is, interventions that an expert physician would not have carried out or would not have carried out at the given time
- derive the results of a variety of laboratory tests directly from the parameters of the "body" of the VP, and
- make the intelligent agent ever more "human-like" by facilitating the modeling of interoception – the conscious perception of symptoms by the model of VP's "mind"
- easily modify the model underlying the simulation as medical discoveries are made with no need to "rewrite" a patient, as would need to be done for narrative-oriented interactive scenarios.

The physiological side of the VP is modeled as a set of interconnected ontological objects representing human anatomy and ontological events representing human physiology, including pathology. Each object is described by a set of ontological properties and their associated value sets. Crucial among the properties are those that link the objects to typical events in which they participate. These events are usually complex – that is, they include other, possibly also complex, events as their components. Following Schank and Abelson (1977), we call these complex events scripts. Examples are swallowing by a healthy VP, swallowing by a VP with a diseased esophagus, the process and effects of interventions, clinically appropriate ways to manage patients with given signs and symptoms, etc.

At any given time, the model of the normal human contains whatever normal anatomical and physiological knowledge was compiled to cover the diseases currently available in the system. So, although at present our virtual humans do not have a highly developed model of the circulatory system, as soon as we have completed the circulatory model – which is currently under development to support the modeling of heart disease – all virtual humans will be endowed with all the associated functionalities and property values.

In MVP, diseases are modeled as processes that cause changes in key property values of a patient over time. For each disease, a set number of conceptual stages is established and typical values or ranges of values for each property are associated with each stage. Relevant property values at the start or end of each stage are recorded explicitly, while values for times between stage boundaries are interpolated. The interpolation currently uses a linear function, though other functions could as easily be employed.

A disease model includes a combination of fixed and variable features. For example, although the number of stages for a given disease is fixed, the duration of each stage is variable. Similarly, although the values for some physiological properties undergo fixed changes across patients, the values for other physiological properties are variable across patients, within a specified range. The combination of fixed and variable features represents, we believe, the golden mean for disease modeling. On the one hand, each disease model is sufficiently constrained so that patients suffering from the disease show appropriate physiological manifestations of it. On the other hand, each disease model is sufficiently flexible to permit individual patients to differ in clinically relevant ways, as selected by patient authors.

At the time of physiological model development, knowledge engineers help physicians to distill their extensive and tightly coupled physiological and clinical knowledge into the most relevant subset, express it in concrete terms, and hypothesize about unknown and currently unknowable things, like the nature of disease progression at the pre-clinical (pre-symptomatic) stage. The models thus created are very different from anything physicians themselves have ever encountered and anything that is available in the medical literature.

The Cognitive Agent

The cognitive agent is endowed with capabilities for perception, reasoning and action. In what follows we briefly describe the features present in the current version of the MVP system.[2]

The cognitive agent is at present capable of two types of perception: interoception and the perception of linguistic input from the human user. Responses to interoceptive input are remembering a sensation (typically a symptom) and deciding whether or not to do anything about it at the given time. Responses to language input include learning (augmenting the agent's ontology, whenever appropriate, as a result of understanding user input), responding to a question or suggestion, generating a question based on information or advice just provided, and remembering the content of dialogs.

MVP uses agenda-style control and goal- and plan-based simulation. The underlying organization of and knowledge representation for the physiological and cognitive agents of the VP are the same. We represent goals as ontological instances of properties. MVP relies on large knowledge bases that include a language-independent ontology of over 10,000 finely described concepts, including the scripts described above; a lexicon, mapped to the ontology, of over 30,000 senses that are expanded at runtime to several times that using productive word-formation rules; a "memory" for each VP that can be populated on the fly during a simulation; and a suite of processors for language processing, reasoning, simulation, and so on.

A core architectural aspect of the system is that all interoception and all language input are automatically translated into expressions in the same metalanguage of memory and reasoning used by the VP. As a result, all cognitive processes are modeled using formal, unambiguous knowledge structures that are grounded in the ontological model of the world. Each VP instance has its own ontology, lexicon and memory, reflecting its own cognitive reality, and each can learn over time through the simulation. For example, the VP can learn information about its disease, possible management strategies, etc., from the user, thus expanding – or changing, if there were previous misconceptions – its ontology and lexicon over time. The state of the VP's knowledge affects all of its decision-making functions.

The VP contains a set of beliefs about each of the other agents in the system that is formulated in exactly the same way as its knowledge about itself – though its knowledge about other agents might be much more sparsely populated. The VP's "profiles" of other agents include their ontologies, fact repositories and lexicons as well as its assessment of the character traits, physical and mental states of these agents as well as their perceived attitudes toward various entities in the world, centrally including the VP itself ("this secretary is rude to me"). Economy in recording this information is a separate issue: for example, the VP may assume that some other agent's ontology and lexicon are almost identical to its own ontology and lexicon and just mark the few differences.

The VP also has a set of beliefs about self that can be different from reality: for example, it can believe itself to be, say, more truthful or more healthy than it is in reality. Some of VP's actions will, therefore, be guided by its real self and some, by its opinion of self. This capability is not yet implemented but is an interesting capability to model in a virtual human, for example, highlighting discrepancies between its behavior and its beliefs about self. The "real" profile of the VP is visible only to system developers but may be revealed to the human user if a teacher considers this pedagogically beneficial,

VPs in MVP vary with respect to the extent of their knowledge of medicine, the content and accuracy of their memory of past events and their personality traits. This enables variation across VP instances in decision making and the content and form of dialog acts which, in turn, enhances the verisimilitude of training encounters. VPs can initiate actions, which also enhances the verisimilitude of training encounters by forcing the user to react to a patient's potentially unexpected behavior. VPs are capable of introspection in that they can explain why they are doing what they are doing. VPs are also capable of learning. As a result of their communication with the user, they can learn facts, ontological concepts with their properties and lexicon entries. These learning-inducing dialogs can be initiated either by the user or by the VP. In the latter case, the process originates with instantiating an instance of the goal "know-information-about-concept" and following the plan that involves asking the user informational questions. In the future, learning by reading and learning by reasoning will be added to the repertoire of the VP.

[2] This discussion of the agent's functionalities is rather coarse-grained. Some more detail can be found in (Nirenburg et al. 2008a,b; McShane and Nirenburg 2009).

Perception

Interoception is the process through which the cognitive agent becomes aware of some of the states of the physiological agent. The result of this awareness is recognizing and remembering a symptom. The communication channel between the physiological and the cognitive agent is narrow: the cognitive agent is not fully aware of the activities of its "body," the physiological agent. In the current version of MVP, the interoception module operates a set of demons that are programmed a) to notice changes in values of specific physiological parameters and b) if these values move outside a certain range, to instantiate corresponding symptoms in the VP's memory – more specifically, in its short-term memory component.

Symptoms are represented as values of properties in the cognitive agent's model of self. This model is realized as an instance of the ontological concept HUMAN and is stored in the cognitive agent's fact repository. In addition to the model of self, the VP also has models of other agents of which the VP is aware, such as each of the MVP users. Note that the properties in this cognitive profile are different from the properties used in the physiological agent. For example, the latter will include the property "pressure in the lower esophageal sphincter" while the former may record a certain level of the property "heartburn" without necessarily knowing that heartburn may be affected by a certain level of pressure in the lower esophageal sphincter.

When certain symptoms appear and reach a certain level of severity, the health-attribute property in the cognitive agent's profile of self decreases. This attribute represents a generalization over all symptoms. A decrease in the value of the health-attribute causes the appearance on the agent's agenda of an instance of the goal "BE-HEALTHY."

Perception through understanding of natural language utterances is a major contributor to enhancing the verisimilitude of a VP encounter. The language understanding module extracts semantic, pragmatic and discourse-oriented meanings from inputs and represents them in the same formalism that is used by the reasoning module. This functionality is useful a) because it facilitates reasoning on the basis of language input and b) because it supports reasoning that is necessary **for** understanding language input in the first place.

Language understanding in the MVP system is based on the OntoSem text analyzer, which integrates a large collection of algorithms and knowledge resources. The list of algorithms used in the current version of the MVP just for basic semantic analysis includes:

- an ambiguity resolution algorithm based on multiple-value selectional restrictions encoded in the OntoSem lexicon and ontology
- a statistically trained word sense disambiguation algorithm based on topic-oriented constraints
- an algorithm for dynamic tightening and relaxation of constraints for cases of residual ambiguity and zero residual output candidates, respectively
- a disambiguation algorithm based on estimating the similarity between pairs of concepts in the ontology by computing weighted distances in ontological space
- an algorithm for the unidirectional application of selectional restrictions to process "unexpected" input words that are not in the system's lexicon
- an algorithm for deriving TMRs for fragments of input in cases of failure to produce TMRs for a complete sentence.

Knowledge resources used in language understanding include grammars, lexicons, an ontology, a fact repository, microtheories of specific language phenomena and various specialized rule sets. The major stages in analyzing a dialog turn are

- Preprocessing – text segmentation, processing of named entities, dates, numbers, abbreviations, punctuation, etc., and morphological analysis
- Syntactic analysis – both constituent structure and dependency structure extraction
- Basic semantic analysis – establishing "who did what to whom," word sense disambiguation, semantic dependency determination
- Discourse analysis – concentrating at present on processing reference (including ellipsis) and speech acts.

The operation of OntoSem is computationally expensive. In an attempt to make the cognitive agent's language understanding more efficient – but also because this capability is something people use routinely – we attempt, whenever possible, to use analogical reasoning to obtain text meaning representations. Analogical reasoning is used for conversational formulas and frequent or common dialog turns. This latter capability is enabled in the MVP by the availability of the memory of assertions (the fact repository).

Reasoning

As mentioned above, reasoning – using knowledge to make decisions – is a component of perception. Outside perception, the following reasoning processes (mental actions) are currently employed in MVP:

- augmenting the memory (fact repository) with the new facts
- learning new lexicon entries and ontological concepts
- augmenting/modifying beliefs about the knowledge and intentions of interlocutors
- making decisions about what to do next:
    - running preference functions to select which goal and plan to pursue next or

- finding analogous situations in hopes of using results of earlier reasoning
- deriving the content to be generated into language.

For example, the goal BE-HEALTHY in the current version of the system has just two plans associated with it: SEE-MD and DO-NOTHING (the latter is selected in hopes that the condition will go away without treatment).

The *reasoning* module in MVP concentrates on processing goals. The inventory of goals is at the moment limited to those relevant to the application. Goals are instantiated as a result of perception ("I feel pain") or as a result of agent's own reasoning ("I don't know enough about the procedure recommended by the doctor to choose a course of action"). The choice of goals to pursue at a given time (some parallelism is simulated) depends on the extent and nature of the VP's knowledge about the world, the contents of its memory of past events and a model of its personality traits, beliefs about self and other agents, genetic predispositions and physical and mental states. Different VPs have not only different personality profiles and beliefs but also different ontologies and fact repositories. The VP in the current version of the system is designed to reason by analogy: each goal is associated with a set of known plans. The current version does not include dynamic planning.

Actions

In the current version of MVP he VP's actions include:

- learning new facts and remembering them
- establishing reference relations between newly learned facts and stored facts
- learning new concepts and lexical items – by being told and by reading
- augmenting/modifying the agenda of goals and plans
- generating language (dialog turns)
- [simulated] physical action (e.g., showing up at the office or the lab, taking medicine, ingesting food/drink that affects health states, etc.).

A typical situation in which the agent will learn by being told is as follows. The user may diagnose the VP with a disease about which it knows nothing other that what is diagnosed is a disease. On the one hand, the user may proceed to describe properties of the disease, in English; alternatively, the VP may ask the user questions about various properties of the disease. The VP knows what kinds of questions can be asked based on its ontological knowledge about diseases in general, and it decides whether or not to ask them based on its personality traits, current mental state, etc. During this interaction the VP will have to understand the text of the user's dialog turn, extract from the text meaning representation the filler or fillers of the property or properties in question and fill them in its ontology (or modify the existing filler if there is one). The lexicon entry for the new lexical unit referring to the disease will contain in its semantic structure zone a direct pointer to the newly learned concept (this is the simplest option available).

Comparisons with Other Systems

This section presents a brief overview of some medically-oriented training systems with the purpose of positioning the MVP system on the map of VP-based training and medical simulation efforts.

Cognitive modeling. As this description of MVP has shown, the environment as well as the patient itself model cognitive abilities and therefore can support non-trivial conversation, collaborative decision-making, realistic simulation, and so on. Other available medical training systems do not seek this level of cognitively-oriented functioning. A common type of medical simulation is realized in technical task trainers, which concentrate on developing motor skills and address cognitive skills only inasmuch as they are necessary for the user to understand a specific technical step, like how to insert a needle. A second type of simulation is realized in non-biomechanistic mannikin trainers (e.g., "SimBaby" by Laerdal, Inc., and "The Human Patient Simulator" by Medical Education Technologies, Inc.), which focus on a narrow scope of acute physiological processes. A third type is not really a simulation but rather a branching narrative scenario, organized as a decision tree, that depicts a specific medical case (e.g., MedCases, Inc.). In these, user options are restricted and responses are highly pre-scripted, being delivered through multiple-choice questions. Most importantly, in such systems patient outcomes are fully predetermined by the prefabricated scenario. A more sophisticated type of medical simulation is claimed to be implemented in the Sim-Patient developed by RTI, Inc., where acute traumatic patient scenarios are made available to a user; however, few details about this system are available as the data structures and content are proprietary.

Several advanced scenario-based VP systems have been developed in the European project eVIP (http://www.virtualpatients.eu/). The systems differ as to delivery methods: e.g., the Web-SP system (Zary et al. 2006; Zary 2007) is designed for internet access but the underlying technology is more or less standard throughout. Thus, VPs are realized through collections of decision trees, with decisions expected to be made by human students on the basis of text presented in the prefabricated decision tree nodes that simulates patient examination and test results (both optionally with the help of pictures and photographs). No actual simulation of the patient's physiology, character traits or cognitive abilities is involved. The center of gravity in this project is in the area of interoperability, ease of authoring and curriculum acceptance of VPs (e.g., Ellaway 2009, Poulton 2009).

The strategic decision to use pre-scripted scenarios and multiple-choice responses – and thus to bypass the complexities of physiological simulation, cognitive modeling and involved natural language processing – led VP developers to concentrate on presentation issues. Videos of human actors, advanced graphics, including avatars, the incorporation of off-the-shelf speech recognition and synthesis software and visualization of results of medical tests (such as X-rays and MRIs) have been prominent among the means of strengthening the verisimilitude of the human-computer interactive experience. Indeed, the latest VP environments (e.g., Courteille et al. 2008) show significant progress even when compared to quite recent ones (e.g., Chesher 2004).

Authoring agents. An unintended negative consequence of relying on prefabricated scenarios is the high cost of authoring patient instances, which limits the utility of scenario-oriented VPs. The scenario-oriented VP community is addressing this issue by attempting to share the burden across the development teams. This brings to the fore the issues of interoperability, standardization and efficient dissemination. However, the cost of developing a single training case remains quite high, requiring, according to some estimates, hundreds of person-hours (see Zary et al. 2006). Creating a training case for a VP that operates autonomously on the basis of a simulation of its physiological and pathological processes takes, as our experience in the MVP project has shown, just a few minutes. But the development of the prerequisites for quick patient authoring – a simulator and a model of disease progression for every disease (or, more precisely, group of diseases) – is, of course, an expensive undertaking.

Is simulation needed? There is one notable medical training environment, CIRCSIM (Illinois Institute of Technology; Evens and Michael 2006), that originally relied on physiological simulation of the baroreceptor reflex but, over time, removed the simulation in favor of a handful of set scenarios that fulfilled all teaching requirements. As Michael and Rovick (1996) explain, “The most effective teaching was being generated from the stored correct predictions for each procedure, not from the quantitative outputs generated by the model”. In MVP physiological simulation is a prerequisite for the VP’s ability to reason independently and initiate behavior. Removing this option would seriously compromise the richness of the repertoire of VP actions and, therefore, be detrimental to its verisimilitude level.

Knowledge-based vs. probabilistic simulation. Walton Sumner and his associates (Sumner and Hagen 2006, Sumner et al. 1996, Marek et al. 1996) developed a virtual patient system for the purpose of advancing medical certification procedures of the American Board of Internal Medicine. This system goes beyond the level of multiple-choice questions and relies on simulating a patient with possibly multiple co-existing conditions and allowing the examinee to intervene. The system addresses the issue of automatically creating instances of VPs by starting with a selected disease (or diseases) and probabilistically creating a hypothesized medical history for each patient instance based on knowledge about the incidence of this disease (or diseases) in the population. This emphasis on generating differentiated patient histories is due to the perceived need for providing a secure testing environment in which patient instances are not reused. The probabilistic nature of much of the operation of this system suggests the use of Bayesian networks as the underlying representational mechanism, which adds complexity to both knowledge acquisition and processing. In MVP there is no need for probabilistic reasoning because VP instances are not probabilistic; they are authored by the teachers and represent a specific, not probabilistic set of choices of disease progression and personality traits. If the teacher needs to present a different set of disease progression and intervention outcomes, he or she does this by authoring a different deterministic VP instance. What is not deterministic in the environment is the behavior of the human user, and there is no reason to develop a probabilistic model of this behavior.

Non-VP medical agent systems. One example of the concept of artificial intelligent agents in the medical domain is the “anthropic agency” approach (Amigoni et al. 2003), implemented as a multiagent environment for modeling and regulating physiological phenomena, specifically, the insulin and glucose levels in diabetes patients. This system relies on what the authors call the anthropic agency architecture, “a powerful paradigm to develop control systems for physiological processes shaped as multiagent systems.” (p. 310). The architecture consists of the traditional steps of perception (called “knowledge extraction”, implemented using a team of identical “extractor agents”), reasoning (“decision making”, implemented using a team of identical “decisional agents”) and action (“plan generation” with their team of “actuator agents”). All the agents in this system are what can be called low-level agents, specialized computer programs (in contrast to high-level agents that simulate a broad spectrum of human-like behaviors). The physiological model has a narrowly directed coverage, no cognitive abilities are simulated for the virtual patient and no network of high-level agents simulating human capabilities is introduced. In general, the complexity of the domain knowledge is not the main focus of this work.

Discussion

In our presentation at the symposium we will discuss in additional detail some of the salient features of the MVP cognitive agent and the different roles that the agent can play both within MVP (e.g., an automated tutor) and in other applications (e.g., an automatic medical advisor environment to assist a variety of types of medical personnel in diagnosing and treating patients).

Additional future applications we are considering include:

a) creating, augmenting and maintaining an encyclopedia of “good clinical practices” that:
- can, in principle, be augmented by any licensed physician;
- is simulation-based, which allows “live” observations of variable disease progressions and the efficacy of different treatments of patients with different genetic and cognitive predispositions;
- will serve as an efficient memory and decision aid;
- can form the basis of refresher courses for physicians and help to prepare for board exams;
b) automatic analysis of digital medical records:
- comparing actual disease and treatment paths with those available through MVP;
- discovering and attesting clinical tendencies and practices as yet “hidden” in those records; and
c) using script-based knowledge available in MVP to support the agent actions for the purpose of improving the understanding and, hence, information extraction and other manipulation of medical texts for more traditional NLP applications.

The knowledge, methods and functionalities present and under development in the MVP project facilitate a variety of extensions, for example:

- Population-level modeling, including simulation of response to epidemics and catastrophic events of varying provenance
- Modeling cultural differences
- Diagnosis and treatment of problems in other subject domains
- Organization modeling.

Finally, we fully recognize that the task we are dealing with is very complex and that joining forces with colleagues working on complementary or similar tasks will be highly beneficial. The following is a partial list of issues on which we think collaboration and teaming will be the most fruitful:

- There is no speech input and output in the current version of the MVP system; it would be very useful to join forces with colleagues to integrate advanced **speech** recognition and synthesis modules adjusted for the needs of the environment;
- The perceptual apparatus in the current system is limited; there is a need and an opportunity to integrate with advanced computer **vision** and **haptics** systems;
- Human-aided knowledge acquisition will continue to be important; there is a need and an opportunity to integrate a variety of **machine learning** algorithms and **efficiency-enhancing tools;**
- Human-computer **interfaces** must be constantly improved; there is a need and an opportunity to integrate latest HCI practices and, possibly, game-oriented interfaces.

References

Amigoni, F., Dini, M., Gatti, N., and Somalvico, M. 2003. Anthropic Agency: A Multiagent System for Physiological Properties. *Artificial Intelligence in Medicine* 27: 305-334.

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

Courteille, O., R. Bergin, D. Stockfeld, S. Ponzer and U. Fors. 2008. The use of a virtual patient case in an OSCE-based exam - A pilot study. *Medical Teacher*. e66-e76.

Ellaway, R. 2009. The realities of the virtual patient. Proceedings of the International Conference on Virtual Patients. Krakow, Poland. June.

Evens, M.; and Michael, J. 2006. *One-on-One Tutoring by Humans and Computers*. New Jersey and London: Lawrence Erlbaum and Associates, Publishers.

Huhns, M., Singh, M., Burstein, M., Decker, K., Durfee, E., Finin, T., Gasser, L, Goradia, H., Jennings, N., Lakkaraju, K., Nakashima, H., Parunak, H., Rosenschein, J., Ruvinsky, A., Sukthankar, G., Swarup, S., Sycara, K., Tambe, M., Wagner, T., and Zavala, L. 2005. Research Directions for Service-Oriented Multiagent Systems. *IEEE Internet Computing*, November-December, 65-70.

Marek, V.W., W. Sumner and M. Truszczynski. 1996. Creating evolution scenarios for hybrid systems. *Proceedings of IEEE-SMC Symposium on Control, Optimization and Supervision*, at CESA96, Lille (1996) 512-516.

McShane, M. and S. Nirenburg. 2009. Dialog Modeling Within Intelligent Agent Modeling. Proceedings of the Sixth Workshop on Knowledge and Reasoning in Practical Dialogue Systems at the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17.

Michael, J.; and Rovick, A. 1996. The Once And For All History of HEARTSIM/CIRCSIM. http://www.cs.iit.edu/~circsim/prehistory.html.

Nirenburg, S.; McShane, M.; Beale, S.; and Jarrell, B. 2008a. Adaptivity in a Multi-Agent Clinical Simulation System. In Proceedings of AKRR'08, International and Interdisciplinary Conference on Adaptive Knowledge Representation and Reasoning. Porvoo, Finland, Sept. 17 - 19.

Nirenburg, S., M. McShane and S. Beale. 2008b. A Simulated Physiological/Cognitive “Double Agent”. In Proceedings of AAAI Fall Symposium on Naturally Inspired Cognitive Architectures, Washington, D.C., Nov. 7-9.

Poulton, T. 2009. From Virtual Patients to Virtual Worlds. Proceedings of the International Conference on Virtual Patients. Krakow, Poland. June.

Schank, R. and R. Abelson. 1977. *Scripts, Plans, Goals and Understanding*. New York: Lawrence Erlbaum and Associates.

Sumner, W. and M. Hagen. 2006. Computer Architecture and Process of Patient Generation, Evolution and Simulation for Computer Based Testing System Using Bayesian Networks As a Scripting Language. US Patent 7,024,399.

Sumner, W., V.W. Marek, and M. Truszczynski. 1995. Data Transformations for Patient Simulations. Proceedings of the 19th Annual Symposium for Computer Applications in Medical Care.

Zary, N. 2007. Virtual patients for education, assessment and research: A web-based approach. PhD Dissertation, Karolinska Institute, Stockholm, Sweden.

Zary, N., G. Johnson, J. Boberg and U.G.H. Fors. 2006. Development, implementation and pilot evaluation of a web-based virtual patient case simulation environment – Web-SP. Medical Education, vol. 6, issue 10.