Disease Interaction in Cognitive Simulations for Medical Training

Marjorie McShane, PhD; George Fantry, MD; Stephen Beale, PhD; Sergei Nirenburg, PhD; Bruce Jarrell, MD
University of Maryland Baltimore County
University of Maryland School of Medicine
{marge, sbeale, sergei}@umbc.edu
Gfantry@medicine.umaryland.edu, BJarrell@som.umaryland.edu

Abstract. Maryland Virtual Patient (MVP) is a simulation and tutoring system for training medical personnel in cognitive decision making skills. It is implemented as an agent network using an ontological knowledge substrate. This paper focuses on disease interaction in the MVP environment, including how diseases and/or their treatments can automatically give rise to other diseases, and how disease manifestation can be altered by another concurrent disease.

1. INTRODUCTION

Maryland Virtual Patient (MVP)[1] is a simulation and tutoring environment for training medical personnel in cognitive decision making skills. MVP is implemented as an agent network that includes both human agents (the trainee and, optionally, a human mentor) and artificial agents (the virtual patient[2], the virtual mentor, lab technicians and specialists). Trainees can use the environment to interview virtual patients; 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. The knowledge in the system is modeled using an ontological substrate that supports not only simulation and decision-making but also the natural language processing requirements of inter-agent communication and automatic mentoring [1].

MVP is being developed by a team from the Institute of Language and Information Technologies of the University of Maryland Baltimore County and the University of Maryland School of Medicine. We believe that MVP has a strong potential to contribute to the teaching of clinical medicine by providing trainees with broad, risk-free, clinical decision-making experience. In addition, the formal models of diseases and best clinical practices developed to underpin the simulation – which are available in machine-tractable and human-readable formats – represent a potent complement to the expository descriptions available in textbooks, as they encapsulate the mental models of domain experts.

Most currently available clinical decision-making systems are not grounded in simulations. Instead, they provide trainees with the opportunity to work through decision trees that target key decision points in the process of diagnosing and treating a disease. MVP, by contrast, offers trainees a more open-ended choice space and a richer scope of interactions, which more closely parallel the demands of actual medical practice.

The current MVP implementation covers six esophageal diseases, several of which include a broad range of clinical manifestations. We chose to initially model esophageal disease because the esophagus is a relatively uncomplicated organ and because one of the symptoms of esophageal disease, chest pain, can cause significant diagnostic dilemmas with cardiac disease. The modeling of cardiac disease is currently underway.

2. DISEASE MODELING OVERVIEW

Detailed descriptions of disease modeling in MVP are available in [2] and [3]. Here we provide only the briefest overview sufficient to support the discussion of disease interactions.

---

[1] Patent pending.

[2] The virtual patient is actually a "double agent", modeled as a combination of a physiological agent and a cognitive agent.

As a modeling strategy, diseases are divided into conceptual stages, each stage being associated with some pertinent event and/or prominent changes of property values in the patient. Property values that change progressively over time are simulated by interpolating between anchor values explicitly associated with each stage.

Consider, for example, gastroesophageal reflux disease (GERD), which will be discussed throughout the paper. (For a more comprehensive description of GERD modeling, see [2].) GERD is modeled as permitting two distinct paths, the first having a maximum of five stages and the second a maximum of three stages, as shown in Table 1. An important clinical aspect of GERD is that not all patients, even if left untreated, will experience all the stages of a given path: for example, an untreated patient might get erosions but never ulcers or a peptic stricture. The GERD model accounts for this variation by permitting patient authors (teachers or specialists who create patient instances from the basic model using a multiple-choice questionnaire) to select inherent GERD-related predispositions for their patients. As such, a patient with a predisposition to erosion but not ulcer will follow path 1 of GERD, but his disease will stop progressing before the ulcer stage, whereas a patient with a predisposition to adenocarcinoma will follow path 2 to its most severe manifestation.

Table 1. Paths and Stages of GERD

	Path 1	**Path 2**
Stage 1	preclinical	preclinical
Stage 2	inflammation	inflammation
Stage 3	erosion	Barrett's esophagus
Stage 4	ulcer	adenocarcinoma
Stage 5	peptic stricture	

Diseases are modeled using a combination of causal chains, for known biomechanical functions, and clinically derived "bridges". The latter are used for (a) poorly understood biomechanical functions and (b) biomechanisms whose details represent a grain size of description not required by the MVP application. Clinical bridges typically associate a change in property values with a period of elapsed time.

GERD, for example, is modeled largely through causal chains: a specific level of exposure of the esophageal mucosa to acidic gastric contents causes inflammation, which, in path 1 patients, causes cell death. If the rate of cell death exceeds the rate of healing, the progressive diminution of cells can cause erosions in the mucosa and ulcers into the submucosa. If, however, the acidity of the reflux is sufficiently reduced – as by effective medications – the rate of cell growth exceeds that of cell death and healing occurs. Within this primarily causal model, bridges permit us to exclude details that are needed neither by the simulator nor by the virtual mentor, like the specific processes that comprise cell death and regeneration. If, later on, we should incorporate into the system substances or processes that can directly affect some aspect of cell death or regeneration, and if those effects could have clinical relevance for the diagnosis and treatment of GERD, then the current bridge will need to be expanded into a full causal chain.

The sections to follow discuss two different types of disease interactions: cases in which diseases manifest themselves differently in combination than they would in isolation, and cases in which one disease or intervention causes another disease. It bears emphasizing from the outset that the goal of this work is to *model* diseases sufficiently to support a robust, realistic interactive simulation, not to attempt to reproduce the human organism from first principles.

3. ONE DISEASE AFFECTS THE CLINICAL MANIFESTATION OF ANOTHER

In some cases, concurrently running the simulations for multiple diseases does not produce the full spectrum of clinically observed manifestations of those diseases in combination. A case in point is LERD-GERD, the concurrence of GERD, a reflux disease affecting the distal esophagus, and LERD (laryngopharyngeal extra-esophageal reflux disease), a reflux disease affecting the proximal esophagus and larynx.[3]

Excessive reflux derives from one of two abnormalities of the lower esophageal sphincter (LES): a consistently hypotensive LES or excessive transient relaxations of the LES (TLESRs). The more reflux per day with increased esophageal acid exposure time, the faster these diseases progress.

[3] Reflux diseases are clinical conditions caused by excessive refluxing of the acidic contents of the stomach or duodenum into the esophagus.

All humans in the MVP system by default have a predisposition to GERD, meaning that if the preconditions for GERD are met, disease processes will initiate. Not so for LERD: in order for a patient to get LERD, an author must specifically assign him this predisposition. As such, GERD patients can be generated "unexpectedly" during a simulation, whereas LERD patients cannot be. This modeling decision reflects the clinical observation that GERD is more prevalent in the population than LERD.

The models for GERD and LERD combine aspects that are fixed across all patients and aspects that can be parameterized to introduce variation among patients. Table 2 shows a subset of the property value correspondences that are fixed for patients with reflux diseases. There are internal correspondences for GERD (left) and LERD (right) as well as correspondences across the diseases. The grayed out area in the GERD table indicates a non-disease state; the remainder of the table shows disease states.

Table 2. Key GERD-LERD property values.

GERD			LERD		
TTAR Distal (% day)	GERD Level	Stage Duration (days)	TTAR Proximal (%/day)	LERD Level	Stage Duration (days)
3	11	-	1.31	9	180
4	11	-	1.4	8	160
5	10	-	1.5	7	140
6	9	180	2	6	120
8	8	160	2.5	5	110
9.5	7	140	3	4	90
11	6	120	3.5	3	60
13	5	110	4	2	50
15	4	90	4.5	1	40
17	3	60	5	0	30
19	2	50			
22	1	40			
25	0	30			

TTAR indicates total time in acid reflux as a percentage of time per day. Since the distal esophagus is anatomically more suited to enduring acid exposure than the proximal esophagus, it takes in excess of 5% (1.2 hours) of daily distal acid exposure to initiate GERD, whereas it takes only 1.3% (30 min.) of daily proximal acid exposure to initiate LERD. Since not all reflux in a LERD patient is expected reach the distal esophagus, TTAR-proximal is a fraction of TTAR-distal: e.g., TTAR-distal 8% corresponds to TTAR-proximal 2.5%.

TTAR determines the values of several properties, including GERD and LERD Level and Stage Duration.

GERD and LERD level are abstract properties whose values are used as a shorthand to refer to a correspondence between a given TTAR, Stage Duration, and what is known as the DeMeester score -- a test result used in the diagnosis of GERD (the latter is included only for completeness' sake, being extraneous to the current discussion). GERD and LERD level indicate the relative severity of GERD or LERD, with lower numbers reflecting a more severe disease state. For patients whose GERD is due to a hypotensive LES, GERD level equals LESP – a useful mnemonic for people interpreting the model.

Stage duration is the duration of each of the conceptual stages of the disease. For example, a GERD patient who has an inherent predisposition to ulcers will experience 4 conceptual stages of path 1 of GERD: preclinical, inflammation, erosion and ulcer (cf. Table 1). If that patient has a GERD level of 7, the duration of each stage will be 140 days, meaning that, if untreated, his disease will reach its most advanced state 560 days after its inception.

Following just the predictions made by Table 2, there should be three categories of reflux patients:

1. Patients who experience GERD but not LERD. For such patients, the value for the inherent property "predisposition to LERD" is "no" (i.e., the ontological default was not manually overridden) and LERD is blocked.

2. Patients who experience LERD but not GERD. In the simplest case, such patients have a TTAR-Proximal of 1.31 – 1.5, which is sufficient to cause LERD, but (reading Table 2 horizontally) a corresponding TTAR-Distal of only 3 - 5, which is not sufficient to cause GERD.

3. Patients who experience LERD and GERD, meaning that they have an inherent predisposition to LERD, a TTAR-Distal of 6% or greater, and a corresponding TTAR-Proximal of 2% or greater.

The above scenarios neatly cover the three obvious classes of reflux patients, however they do not account for all clinically observed manifestations of LERD(-GERD). Specifically, in the presence of LERD, GERD can show a physiological and/or symptom-related manifestation that deviates from the expectations encoded in the basic GERD model [2].

- Physiological deviation: Some LERD(-GERD) patients experience a slower rate of disease progression than would be

predicted by their GERD level. Such patients might have highly effective mucosal defense against distal esophageal acid-induced damage.

- Symptom-related deviation: Some LERD(-GERD) patients experience lower than expected levels of symptoms given, given their GERD-related physiological state. Such patients have decreased visceral sensitivity to distal esophageal acid exposure.

The question arises: should the basic GERD model be expanded to cover these irregularities, or should unexpected manifestations of GERD in the presence of LERD be handled separately? Our answer is pragmatic: since the basic GERD model functions robustly in the absence of LERD, LERD-associated anomalies of GERD are handled through disease interaction parameterization.

Two properties were introduced to scope over the manifestation of GERD in the presence of LERD:

1. **GERD physiological manifestation**, which indicates how the typical duration of GERD stages (as determined by the independent GERD model) will be affected by GERD's co-occurrence with LERD. This property is not intended to imply that GERD progression is actually, physiologically, affected by LERD; rather, it is used to model the clinical observation that – for reasons unknown – some LERD patients who would be expected to have GERD due to their level of distal acid exposure either do not experience the related physiological changes or experience them at a slower rate of progression. Values for this property range from 1 to 200, with 1 meaning that GERD progresses at its normal rate and 200 meaning that it progresses 200 times as slowly as would be expected, effectively never reaching the clinical stage of the disease (the latter would model the case of a person having an "iron distal esophagus").

2. **GERD sensitivity**, which indicates how sensitive a LERD patient is to the symptoms of GERD. This property concerns only the subjective experiencing of symptoms, which can change dramatically in the face of other concurrent symptoms (e.g., one temporarily forgets about a low-grade headache upon stubbing one's toe). The values of this property range between {0,1}, with 1 meaning that GERD symptoms will be experienced in full, .5 meaning that they will be reduced by half, and so on.

Including these variables into the process of authoring LERD(-GERD) patients adds two more classes of options to the original patient parameterization provided by Table 2, such that the scope of profiles of LERD(-GERD) patients in MVP more closely matches a real clinical population.

A natural question would be, why not include one or both of these variables in the process of authoring GERD as an independent disease, thus permitting a broader scope of GERD patients in the resulting population? After all, one would suspect that certain individuals who have no predisposition to LERD could experience unexpectedly mild physiological or symptom-oriented manifestations of distal acid exposure. There are at least two arguments in favor of not relaxing the core model of GERD. First, doing so would dilute the model's predictive power and potentially mask the core generalizations the system seeks to convey to trainees. In addition, there is no clinical evidence about patients who have no esophageal symptoms: they may or may not be experiencing levels of acid exposure that would be expected to cause symptoms. The only reason we know about some "unusual" GERD patients is because they also have LERD and undergo testing for that disorder. In sum, taking all constraints off of the GERD model would not serve the educational goals of MVP. That being said, when we reach the stage of configuring MVP for highly advanced trainees, we might want to include the potential to create cases that stray from the basic model, to emphasize that any clinically derived mode is subject to exceptions.

4. A DISEASE OR TREATMENT SPAWNS ANOTHER DISEASE

As described above, GERD is a disease that can be automatically generated in patients who are not explicitly authored to have GERD. The reason it can be automatically generated is due to the causal chain that underlies the GERD model: GERD is initiated whenever a patient's LESP drops below 10 mmHg, regardless of the cause of that hypotensive state. In this section, we describe a number of internal and external influences that affect LESP as well as the evaluation function that determines LESP in the context of multiple simultaneous affecting factors.

At any point in the simulation, a patient's LESP is a function of the LESP at the previous simulation point and the effects of any of the

three types of events listed in Table 3, which are described below in turn.

Table 3. Agents and Events that Affect LESP

Type of agent	Event
Physiological agent of VP	Progression of diseases that affect LESP
Cognitive agent of VP	Reflux-related lifestyle choices
Trainee and specialist agent	Interventions that affect LESP

Progression of diseases that affect LESP. Currently, MVP includes two diseases that directly affect LESP: scleroderma esophagus reduces LESP over time whereas achalasia increases it. The value of LESP as determined by disease progression is interpolated between so-called "clock cycles" of the simulation – essentially, points at which the simulator recalculates property values in response to the passing of time or internal or external stimuli. It was decided that no patient in MVP will simultaneously have scleroderma esophagus and achalasia since this is an extremely rare if at all possible combination of diseases about which there exists no body of clinical knowledge.

Lifestyle choices. Some patients with GERD have food sensitivities: ingesting such foods lowers LESP and/or increases the number of daily TLESRs. Rather than attempt to model each instance of ingesting such edibles and tracing their short- and long-term effects on LESP (which would likely be impossible and certainly is unneeded for MVP simulation), we employ the following generalization: on days when a patient ingests an offending substance, his GERD level is 1 less than for days when he is "clean". For example, a caffeine-sensitive patient whose GERD level is 8 on caffeine-free days will have a GERD level of 7 on days when he succumbs to temptation (cf. related effects in Table 2). This relatively small change in GERD level reflects that clinical observation that while changing food-related habits can have some effect on GERD, only in the most mild GERD patients are lifestyle modifications sufficient to reverse the course of disease. From the point of view of simulation, a lifestyle modification – for good or bad – changes the calculation of LESP only on the first day it is in effect. For example, if caffeine-sensitive GERD patient stops drinking coffee his GERD level will increase by 1 and remain increased until he goes back on caffeine. A GERD patient's lifestyle can change as frequently as every day.[4].

Interventions that affect LESP. Three interventions can directly affect LESP: Heller myotomy, pneumatic dilation and BoTox injection. All of these are used to treat the esophageal disease achalasia and all of them, if successful , decrease LESP.

- Heller myotomy decreases LESP by surgically cutting the LES. It can have three outcomes: (a) long-term success, meaning that LESP reaches a level of 0 – 10 mmHg and remains low over time, no longer subject to the progressive tightening caused by achalasia; (b) success with regression, meaning that LESP is initially low but is again subject to the gradual tightening caused by achalasia; (c) failure, in which the LESP remains hypertensive.
- Pneumatic dilation is an endoscopic procedure that uses an inflated balloon to tear the LES, thereby decreasing its pressure. The outcomes of pneumatic dilation parallel those of Heller myotomy, with the only difference that the resulting LESP tends to be a bit higher – between 4 and 12 mmHg.
- BoTox is a neurotoxin derived from the bacteria, *Clostridia botulinum*. When injected into muscle, the toxin inhibits release of acetylcholine from presynaptic neurons preventing the nerve impulse from reaching the muscle, which results in muscle relaxation or paralysis. The net effect in achalasia is to improve the balance between inhibitory and stimulatory innervation of the LES resulting in a decrease in its basal or resting tone. BoTox injection tends to decrease LESP less significantly than the other two procedures and its effects always diminish over time.

All three of these treatments are considered good practice only for patients with achalasia. Accordingly, it is only during the authoring of achalasia patients that a physican-educator is asked to explicitly select how his patient will respond to each treatment if the trainee should administer it. For example, a patient author might decide that if his achalasia patient is

[4] The patient's decisions regarding lifestyle are currently modeled rather simply, but work has begun on expanding the cognitive side of the virtual patient to permit a closer modeling of lifestyle choices.

given a pneumatic dilation, the procedure will be reduce LESP to 4 mmHg with regression to 50 mmHg over 5 years, whereas if the patient is given a Heller myotomy, it will reduce LESP to 1 with no regression. However, what if one of these procedures is performed on a patient who does not have achalasia and who is not explicitly assigned an outcome for these procedures? In that case, ontological defaults are applied, as we will now describe.

Every virtual patient, as a virtual human, has values for every property in the system. In principle, patient authors could be asked to explicitly select every property value for every patient, but this would be both time consuming and unnecessary. Looking a few years down the line, this would mean asking the author of a patient with a torn anterior cruciate ligament to indicate how his patient would respond if given a Heller myotomy! However, although authors should not be burdened with such parameterization, the simulator must know how to respond no matter what a trainee does, whether it be launching a Heller myotomy on a knee patient or – for reasons unknown – asking a patient what his eye color is.

In order to prepare the simulator for every eventuality, every patient instance is provided with a full inventory of property values, many of which are assigned by default or random selection. For example, if a Heller myotomy is performed on a patient who does *not* have achalasia, his LESP will reduce to 2 mmHg and remain at that level permanently. This is an ontological default assigned to any patient who is not authored to have a specific response to a Heller myotomy.

At this point, the mechanisms for automatically generating GERD should be clear.

- If a patient has scleroderma esophagus, his LESP will eventually reduce to a level that initiates GERD.
- If a patient with level 10 GERD (a non-disease state; cf. Table 2) engages in bad dietary habits, his TTAR will increase (reflected as a GERD level of 9), which will initiate GERD.
- If a trainee performs a Heller myotomy, pneumatic dilation or BoTox injection with a resulting LESP of < 10 mmHg, GERD will be initiated.

When GERD is automatically generated rather than authored, its path is determined by the GERD-related property values that were assigned to the patient upon its generation. (Recall that all patients must have values for all ontological properties, with property values not relevant to the specific disease being assigned randomly or by default.) For example, the GERD path that is currently used as the default is path 1, stopping at ulcer. In the future we plan to introduce more sophisticated automatic property value selection functions into MVP. Of course, patient authors can explicitly select whichever properties they want for their patients, expanding beyond the narrow subset required to insert the abnormal pathology related to one or another disease into a patient.

An evaluation function recalculates the LESP after any clock cycle during which any of the LESP-affecting factors have fired (cf. Table 3). Consider an achalasia patient whose LESP would have risen to 45 mmHg by a given point in time, but at that time he is given a successful Heller myotomy: the effect of the Heller myotomy “trumps” the LESP effect that that disease would have had, and the patient’s resulting LESP will be that determined by the Heller myotomy.

5. EVALUATION

Recently, a prototype of the MVP simulation and tutoring environment was used with forty second and third year medical students at the University of Maryland School of Medicine to assess the functionality of the MVP. Through student interaction with the MVP, a number of key MVP capabilities were demonstrated and a number of observations were made.

The MVP realistically presented esophageal diseases that evolved over weeks, months and years including symptoms, diagnostic findings and responses to therapy which were uniformly congruent. In addition, the MVP autonomously developed a complication of a treatment, namely GERD developing after a Heller myotomy. Importantly, students were able to observe the natural history of the disease, effectiveness of treatment and recurrence of symptoms over time. This is a unique learning opportunity as opposed to seeing a patient with a disease at a single point in time during a clinical rotation.

While evaluating and treating the MVP, students were able to view the pertinent medical record including the narrative history, the results of diagnostic tests and response to treatment as well as the underlying physiology simultaneously. Viewing these three aspects simultaneously with the associated physiology serves to reinforce underlying pathophysiology associated with the disease as well as allowing the student to observe pertinent

changes with disease advancement or as a response to treatment.

When assessing student performance, several important observations were made. With the MVP behaving and evolving autonomously in a very realistic manner, the students rapidly became engaged in managing the patient as they were faced with clinical dilemmas with consequences to their actions. One of the more challenging issues that arose on several occasions is whether to observe or take further action, such as ordering an additional diagnostic test or prescribing additional treatment. This was difficult for the students and identified a skill requiring further refinement.

Students presented with initial complaints of a patient were able to evaluate and manage the patient in a manner that parallels the actual process in clinical care. Students could query the patient, order diagnostic tests, or prescribe treatments. After advancing time, students were able to observe the results of therapeutic interventions, reassess the patient's symptoms, and make further management decisions as warranted. This allowed the students to observe and manage disease at onset and over extended periods of time as with a chronic disease.

Second year students were able to apply their basic science knowledge to the clinical scenarios and third year students were able to manage patients in a similar manner to their clinical rotations. Students demonstrated longitudinal thinking as required when managing a patient over time, addressing such questions as "When should the patient return for a follow-up visit?" and "How long should it take for the treatment to have a benefit?" Furthermore, students were faced with clinical uncertainty which initially was a difficult concept for them. These are clinical challenges not easily reproduced in a single clinical setting or on a single best answer multiple choice question.

In addition to observing the natural history of a disease and the effects of treatment over extended periods of time, students were able to manage a variety of patients including patients with similar symptoms but different underlying diseases (surface similarity) and patients with different presentations of the same disease (depth similarity).

Students selectively utilized the tutoring environment. With the tutor engaged, a hypothesis or working diagnosis is required for ordering a test or prescribing a treatment. Often, students did not consciously provide a preliminary diagnosis prior to moving ahead with attempted diagnostic tests and management. The tutor was very effective in this setting, redirecting the students to synthesize their thoughts based on available data related to a specific patient and to establish a working diagnosis.

In a preliminary evaluation of the MVP simulation, students viewed the MVP as a potentially valuable educational tool that simulated clinical scenarios in a more authentic manner that cases typically presented in lecture or small groups. Students felt that the MVP could either be used in small group format with a preceptor or as a tool for self-study.

6. RELATED WORK AND FUTURE DIRECTIONS

Building knowledge-based systems for large domains is a labor-intensive undertaking. As such, most development efforts either apply knowledge-based methods to a relatively narrow/simple domain (e.g., cargo shipping in TRAINS/TRIPS [4] or providing information about movies and restaurants in DUDE [5]) or they pursue a broad/complex domain but using stochastic approaches (e.g., there are medical testing systems [6] and medical decision making trainers [7] that rely on Bayesian networks). However, stochastic approaches have one drawback when used for pedagogical systems: a lack of predictive power about what, exactly, the system will do, and the inability to observe – and therefore explain – how and why the system produced the given output. MVP attacks the problem of domain complexity using a novel two-pronged strategy: (a) encoding detailed knowledge about the nature of objects and processes in a formal ontology, which supports both its organization and its reuse, and (b) selecting a grain size of description to suit each particular phenomenon covered in the application.

As regards tutoring systems, among the most established in the field of medicine is CIRCSIM-Tutor, whose focus has fundamentally shifted since its inception [8]. Initially, under the name MacMan, the system was a mathematical model of the baroreceptor reflex that could be explored by students but provided no feedback. Over time, the mathematical model was removed (being replaced by stored correct predictions) and the dynamic aspects were constrained to the tutoring process itself. In MVP, by contrast, the

simulation is and will remain central to achieving the pedagogical goals.

The current prototype of the MVP system showcases the simulation of the physiological functioning of the virtual patient – that is, the physical side of this physiological + cognitive "double agent." The next stage of work, which we are just beginning, involves adding cognitive function to all human-like system agents: the cognitive side of the virtual patient, the virtual mentor, specialists, lab technicians, and so on. Cognitive function includes selecting and managing plans and goals, communicating in natural language (the current prototype uses menu-based communication), and displaying the effects of personality traits on disease course and health-care interactions (affective modeling). In order to keep cognitive modeling feasible within MVP, all aspects will be pursued with an application-driven focus. The contributions of D. Traum, e.g., [9], have been particularly useful as we define the problem space for our plan-based dialogue modeling.

We have also recently begun work on heart disease, which offers a particularly interesting testbed for both physiological and cognitive modeling.

REFERENCES

[1] Sergei Nirenburg and Victor Raskin, *Ontological Semantics*. MIT Press, 2004.

[2] Marjorie McShane, Sergei Nirenburg, Stephen Beale, Bruce Jarrell and George Fantry, Knowledge-based modeling and simulation of diseases with highly differentiated clinical manifestations. 11th Conference on Artificial Intelligence in Medicine (AIME 07), Amsterdam, The Netherlands, July 7-11, 2007.

[3] McShane, M., B. Jarrell, S. Nirenburg, G. Fantry, S. Beale. Forthcoming. Training Clinical Decision Making Using Cognitively Modeled Virtual Patients. In: Park A, Witzke D, Klein R eds., *Minimally Invasive Surgery Training: An Electronic Book*. In Press: University of Maryland School of Medicine.

[4] George Ferguson and James F. Allen, TRIPS: An Integrated Intelligent Problem-Solving Assistant, *Proceedings of the Fifteenth National Conference on AI (AAAI-98)*, Madison, WI, 26--30 July, 1998.

[5] Oliver Lemon and Xingkun Liu, DUDE: a Dialogue and Understanding Development Environment, mapping Business Process Models to Information State Update dialogue systems. *Proceedings of EACL 2006*.

[6] W. Sumner 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.

[7] Susan McRoy, Syed S. Ali and Susan M. Haller, Uniform knowledge representation for language processing in the B2 system. *Journal of Natural Language Engineering*, 3(3), 1997.

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

[9] David Traum, *A Computational Theory of Grounding in Natural Language Conversation*. PhD Thesis and TR 545, Computer Science Dept., U. Rochester, December 1994.