

Revealing the Conceptual Substrate of Biomedical Cognitive Models to the Wider Community

Marjorie MCSHANE^a, PhD, Bruce JARRELL^b, MD,
George FANTRY^b, MD, Sergei NIRENBURG^a, PhD,
Stephen BEALE^a, PhD, Ben JOHNSON^a

^a *University of Maryland Baltimore County*

^b *University of Maryland School of Medicine*

Abstract. The patient authoring interface for each disease in the Maryland Virtual Patient simulation environment reveals the conceptual substrate of the disease model. Revealing the disease model to the community both explains how the interactive simulations work and invites collaboration from the wider community.

Keywords. cognitive simulation, virtual patient, medical education

Background/Problem

Many biomedical simulations are designed such that models of causal or temporal reasoning are either not used or are not made inspectable. However, in certain modeling and simulation domains, it is useful for the underlying models to be inspectable by the broader community, both for evaluation and for collaboration. The simulations in the Maryland Virtual Patient (MVP) Project are based upon ontologically encoded cognitive models of diseases that reflect the mental models of practicing physicians. These models are available in encapsulated form in the interface used to create instances of virtual patients – a process that anyone, not only developers, can carry out. As such, the outside community is welcome to evaluate the core models and to suggest modifications, which might reflect alternative opinions or new findings. The ontological organization of knowledge makes it easy to incorporate such modifications into the simulation environment.

The original version of the MVP system was described in [1]. Since that time, we have completed the modeling of several more diseases and have significantly expanded the patient authoring environment, understanding the importance of making our disease models transparent to outside physicians and biomedical researchers.

Tools and Methods

MVP currently supports the simulation of six esophageal diseases, some of which have many possible tracks of clinical manifestation: achalasia, gastroesophageal reflux

disease (GERD), laryngopharyngeal extraesophageal reflux disease (LERD), LERD-GERD (a combination of LERD and GERD), scleroderma esophagus and Zenker's diverticulum. The model for each disease, which is created as a collaboration between physicians and knowledge engineers, is encapsulated in the interface used for authoring instances of patients with that disease.

The patient creation process for all diseases begins with providing basic information about the patient: name, age, gender, weight, race, etc. We omit this aspect of the interfaces, as well as other aspects that are easily described in prose, in the screen shots below for reasons of space.

Achalasia is a disease that progressively renders a patient unable to swallow, which is thought to be due to a loss of relaxing neurons in the lower esophageal sphincter (LES). This disease is modeled as having five stages, t0 through t4, with t0 being preclinical. The duration of each stage is variable across patients.

	t0	t1	t2	t3	t4
Stage Duration (in Months)	12	12	12	12	12

The independent variable in achalasia is the ratio of relaxing (inhibitory) to contracting (stimulatory) motor neurons in the LES, which decreases steadily over the five stages of the disease. All other physiological properties depend upon this variable. However, for purposes of modeling, we refer to the basic LES pressure as a stand-in for independent variable since the correspondence between the neuron ratio and basic LES pressure is constant and it is easier for physicians to orient diagnostics and treatment around LES pressure. The disease model for achalasia contains relatively few parameterizable variables (in orange cells): the duration of each

Physiological Properties	Start	t0	t1	t2	t3	t4
Ratio of relaxing to contracting neurons in the distal esophagus	100/100	80/100	60/100	40/100	20/100	10/100
Basal LES Pressure	25	t0 + 8	t0 + 16	t0 + 24	t0 + 32	t0 + 40
Residual LES pressure (torr)	0	8	16	24	32	40
Residual LES diameter (cm.)	2.0	1.5	1.0	0.5	0.0	0.0
Amplitude of contraction during peristalsis	80	65	40	30	20	10
Efficacy of peristalsis	peristalsis	peristalsis	peristalsis	intermittent	aperistalsis	aperistalsis
Diameter of distal esophagus (cm.)	2	2.8	3.6	4.2	5	6
Retained esophageal content (on the abstract scale {0,1})	0	.1	.3	.55	.85	1
Emptying delay (min.)	0	1	5	10	30	35000

Symptoms	Start	t0	t1	t2	t3	t4
Difficulty Swallowing Distal	0	0.1	1	2	3	4
Do solids stick?	no	no	yes	yes	yes	yes
Do liquids stick?	no	no	no	yes	yes	yes
Weight loss	0	0	0	0	.1	.2
Chest pain	0	0	.1	.3	.5	.7
Regurgitation (times/month)	0	0	0	10	40	70

Figure 1. An excerpt from the authoring interface for achalasia.

Disease models break down into two major classes based on whether or not the physiological causal chains underlying the disease are well understood. In cases where physiological causal chains are relatively poorly understood – as for achalasia, scleroderma esophagus and Zenker's diverticulum – the simulation is primarily driven by temporal causal chains. Each disease is divided into conceptual stages, with each stage being associated with clinically observed physiological changes and symptom profiles. As simulated time passes, the patient's state changes incrementally, calculated using an interpolation function that incorporates the start value of each property at the beginning of the disease and the end value for each conceptual stage. Figure 1 shows these aspects of the model of achalasia, as presented to patient authors. The text in the blue background explains each aspect of the model using methods of progressive

disclosure (a small text field with a scroll bar), which permits users with different levels of experience using the system to use the same interface without the explanatory materials becoming cumbersome. The explanatory text conveys important aspects of how the recorded property values are interpreted within the model and processed by the simulation engine, like which property values impaired by a disease can be reversed given an effective treatment, and which variables are independent and which are dependent. In short, the knowledge used by the simulator goes well beyond what is needed to parameterize a new patient instance, and it is conveyed to patient authors as text in order to clarify – albeit in encapsulated form – how the instantiated model works.

The gray cells indicate values that are fixed for all patients, since permitting their variation is not necessary for either of our immediate goals: (a) generating automatic function in the simulation (e.g., if a given biological pathway can be affected by medication, then it must be parameterizable) and (b) permitting noteworthy variation among patients within a teaching context. The orange cells indicate property values that can be changed for each patient, within ranges visible by mousing over the given cell. This division between parameterizable and non-parameterizable property values points up an important benefit of making our models accessible to the community: for the current teaching application, it was appropriate to make certain property values parameterizable within certain ranges; however, for some other application it might be necessary to make more of these values more variable across patients, which can be readily done with no changes required of the simulation engine.

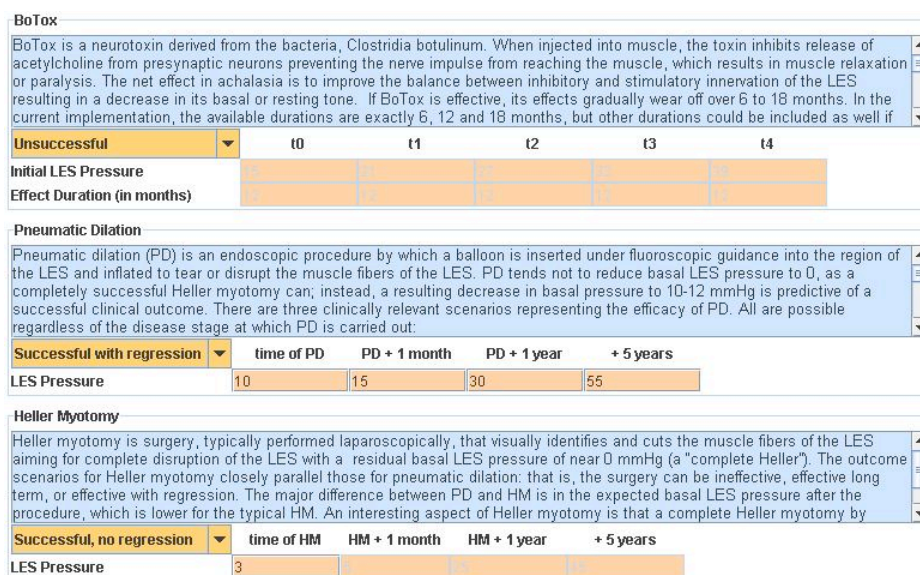


Figure 2. Treatment outcomes for achalasia patients.

The remaining aspect of patient parameterization for achalasia regards treatments. There are three treatment options, each explained in the associated blue shaded text fields (see Figure 2). Each has three potential options: unsuccessful, successful with regression, and successful without regression. If a treatment is unsuccessful, as for BoTox in Figure 2, there are no further choices to be made: the patient's condition is

unchanged (note that the values in the cells of the corresponding table are grayed out). If a treatment is successful with regression (as for pneumatic dilation in Figure 2), the author must choose the rate of regression of the basal pressure of the lower esophageal sphincter (LESP) over time: its value one month, one year and five years after the procedure. (Although LESP is actually dependent upon the ratio of contracting to relaxing neurons, it is conceptually easier for clinicians to reason using LESP). After a procedure, most physiological properties and symptoms retain the original correspondences with LESP shown in the tables in Figure 1; however, the efficacy of peristalsis and diameter of the distal esophagus never improve once compromised, as explained in the blue text field. If a treatment is successful without regression (as for Heller myotomy in Figure 2), only the original post-procedure LESP must be indicated, with most other property values following suit, as described above.

The other class of diseases modeled in the system are those for which physiological causal chains are quite well understood. GERD, LERD and LERD-GERD are all of this type. For reasons of space, we highlight just one aspect of the causal modeling of GERD and how it is reflected in the patient authoring interface (see [2] and [3] for more in-depth descriptions of these disease models).

GERD can be defined as any symptomatic clinical condition that results from the reflux of stomach or duodenal contents into the esophagus. Based on a person's inherent predispositions (no biomarkers have yet been discovered), the disease can take one of six paths, shown at the top of Figure 3. The author selects one path for his patient, which sets associated property values in the patient. The two sources of GERD are abnormally low pressure of the lower esophageal sphincter (LES) (< 10 mmHg), or an abnormally large number or duration of transient relaxations of the LES (TLESRs), both of which result in increased acid exposure of the esophageal lining. The text in blue in Figure 3 (which is quite long; note the slider size) describes how LESP and/or TLESRs are used as independent variables in the model. We repeat an excerpt from that text here as an example of how text complements the formal (in some cases, mathematical) aspects of disease models.

The severity of the GERD-producing factors is reflected by the attribute "GERD level", which was introduced to unify the model, abstracting away from which specific LES-related abnormality gave rise to the disease. The lower the GERD level, the higher the daily esophageal acid exposure and the more fast-progressing the disease. The reason for associating a low GERD level with severe GERD is mnemonic: the GERD levels are the same as the basal LESP for patients who have low-pressure GERD. For example, a patient with a LESP of 1 mmHg will have a GERD level of 1. If a patient has a GERD level of 1 due to TLESRs, that means his daily esophageal acid exposure from the transient relaxations is the same as it would have been if he had had a basal LESP of 1. Using GERD level as the anchor for modeling provides a simple mechanism for incorporating a patient's lifestyle habits into the simulation: whenever he is engaging in bad lifestyle habits (assuming he has GERD-related sensitivities to those habits), his GERD level decreases by 1. For patients with a baseline GERD level of 10 – which is not a disease state – this means that engaging in bad habits is sufficient to initiate GERD and discontinuing them is sufficient to promote healing without the need for medication. For patients with a baseline GERD level of less than 10, lifestyle improvements can slow disease progression but not achieve the healing of previous esophageal damage.

- inflammation
- inflammation > erosion
- inflammation > erosion > ulcer
- inflammation > erosion > ulcer > peptic_stricture
- inflammation > Barrett's_esophagus
- inflammation > Barrett's_esophagus > tumor

The severity of the GERD-producing factors is reflected by the attribute "GERD level", which was introduced to unify the model, abstracting away from which specific LES-related abnormality gave rise to the disease. The lower the GERD level, the higher the daily esophageal acid exposure and the more fast-progressing the disease. The reason for associating a low GERD level with severe GERD is mnemonic: the GERD levels are the same as the basal LES pressure for patients who have low-pressure GERD. For example, a patient with a LES pressure of 1 mmHg will have a GERD level of 1. If a patient has a GERD level of 1 due to TLESRs, that means his daily esophageal acid exposure from the transient relaxations is the same as it

Select one	TTR in hours	TTR in %/day	LES Pressure	Stage Duration
GERD Level 10	1.2	5	10-15	N/A
GERD Level 9	1.56	6.5	9	180 days
GERD Level 8	1.92	8	8	160 days
GERD Level 7	2.28	9.5	7	140 days
GERD Level 6	2.64	11	6	120 days
GERD Level 5	3.12	13	5	110 days
GERD Level 4	3.6	15	4	90 days
GERD Level 3	4.08	17	3	60 days
GERD Level 2	4.56	19	2	50 days
GERD Level 1	5.28	22	1	40 days
GERD Level 0	6	25	0	30 days

Transient LES Relaxations Low LES Pressure Enter the GERD level:

Figure 3. An excerpt from the patient authoring process for GERD.

As is clear by the table, when the author selects the GERD Level (he chose 7 in Figure 3) the duration of each stage of the disease and the total time in reflux (TTR) are automatically selected for him. Other aspects of the patient authoring interface permit authors to select lifestyle habits for their patients, whether those lifestyle habits affect their GERD, their symptom profile and their response to medications. Our point in this example is to show that even when a disease is modeled using causal chains that are encoded in quite complex ontological scripts and realized in even more complex simulation programs, the conceptual substrate of the basic models can readily be shared with – and contributed to – by the larger community.

Discussion

The patient authoring interface in MVP highlights key aspects of the cognitive model of each disease, providing patient authors with explanations of the choice space without either repeating all the information about each disease available in textbooks or expounding upon the implementation of the simulation engine. The core aspects of each disease model include which property values are parameterizable among patients and which ones are fixed for all patients, what ranges of values are permitted for each property at each stage of the disease (seen by rolling over cells in the interface), how “healing” is interpreted with respect to each property value affected by a disease, and how parameterizable property values are used to “bridge” unknown aspects of diseases, like as yet undiscovered genetic influences. The grain-size of description – including which aspects are made parameterizable and which physiological causal chains are included in the model – is influenced by the given application but could easily be changed to suit other applications using our ontologically grounded knowledge encoding methodology. Let us consider this last point in more detail using the example of GERD. As shown in Figure 3, by selecting a GERD level, the author automatically sets the duration of each conceptual stage of the disease and the total time in reflux per day. For a pedagogical application, these fixed correspondences are very useful.

However, we plan to use this knowledge environment for other applications as well, like automatically analyzing electronic patient records both to validate the model and to learn new population-level clinical knowledge. It is likely that some patients fall outside the range of expected outcomes of our current model, which can lead in two directions: either expanding the current model by making more aspects parameterizable, or creating a second, non-pedagogical version of GERD, thus permitting the pedagogical version to retain strong correspondences that are useful as a conceptual architecture, abstracting away from confounding cases. In fact, our knowledge environment can accommodate any number of versions of a disease model suited to different applications. Similarly, the mentoring module (which we did not discuss here) can also accommodate any number of versions: currently, our virtual mentor reflects one set of clinical preferences, but there could be a entire population of virtual mentors reflecting variations on clinical management practices. Recording disease and mentoring models using ontological scripts (rather than, for example, very large rule sets) permits such variation to be readily recorded and managed.

We have designed our knowledge environment such that we can readily collaborate with the broader community. For example, as more genetic influences on diseases are discovered, and more causal chains are understood, these will be used to replace temporal causal chains with physiological ones. We have made our disease models inspectable so that not only can experts assess them in terms of how virtual patients behave in a simulation, but also in terms of the core tenets of the mental models of the physicians who contributed to their development.

References

- [1] B. Jarrell, S. Nirenburg, M. McShane, G. Fantry, S. Beale, D. Mallott and J. Razcek. 2007. An interactive, cognitive simulation of gastroesophageal reflux disease. *Proceedings of Medicine Meets Virtual Reality 15*, Feb. 6-9, 2007, Long Beach, California.
- [2] M. McShane, G. Fantry, S. Beale, S. Nirenburg, B. Jarrell. 2007. Disease Interaction in Cognitive Simulations for Medical Training. *Proceedings of MODSIM World Conference, Medical Track*, 2007, Virginia Beach, Sept. 11-13 2007.
- [3] M. McShane, S. Nirenburg, S. Beale, B. Jarrell and G. Fantry. 2007. 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.