Proceedings of the 1996 Winter Simulation Conference
ed. J. M. Charnes, D. J. Morrice, D. T. Brunner, and J. J. Swain

MODELING DISEASE PROCESSES FOR DRUG DEVELOPMENT: BRIDGING THE GAP BETWEEN QUANTITATIVE AND HEURISTIC MODELS

Pamela K. Fink
L. Tandy Herren

Medical Science Systems, Inc.
100 N.E. Loop 410, Suite 1350
San Antonio, Texas 78216, U.S.A.

ABSTRACT

Drug development is a complex and extremely risky process. It involves developing knowledge of the relevant biology at a variety of levels of detail in order to piece together a coherent model of the disease process and the potential point(s) of intervention. Ultimately, this model must be complete enough to support the prediction of effects and the explanation of clinical trial results. Unfortunately, such models of a disease process have been primarily developed and maintained in human minds. Consequently, such models are limited by a human's ability to store all of the needed information, structure it properly, and reason about its consequences involving both complex feedback and time dependencies. The following paper describes a knowledge-centered approach to the development of biological models that support the drug development process. This technique focuses on a multi-level, hierarchical approach that models various levels of the related biology using appropriate representation techniques based on both the type and availability of the knowledge.

1 THE DRUG DEVELOPMENT PROBLEM

New drug development is an extremely speculative endeavor. In order to bring a new drug to market, numerous hurdles must be overcome. Each hurdle involves gaining knowledge about how the drug works, under what situations it works, and whether or not it is safe. New drug development is typically motivated by a need or opportunity to intervene in a particular disease process to improve a patient's quality of life. This improvement could be based upon alleviation of symptoms, slowing the progression of symptoms, or actually curing the disease. However, though the success of a drug intervention can only be evaluated at the patient symptom level, current breakthroughs in potential drug interventions are at the sub-cellular level. Thus, for the drug development process to be successful (i.e., a drug has been identified, tested, marketed, used, and deemed effective and safe), several very complex linkages must be made:

1) the signs and symptoms of the disease must be linked to the basic biology that causes them,
2) an intervention must be identified that will alter the basic biology such that the disease signs and symptoms are altered in the desired manner, and
3) a chemical compound must be identified that alters the basic biology of a human in the desired way, but without any major undesirable effects.

To understand how a drug acting at the sub-cellular level can affect a disease's symptom complex, the individual cell behavior, the behavior of pools of cells, and the cellular interactions within a specific biological environment must be known. In addition, how cell behaviors and the various chemicals work to affect the symptom complex at the patient level must also be pieced together.

The problem, of course, is that there are a large number of unknowns and uncertainties in how all of the pieces fit together, which makes the process of drug development extremely unpredictable. In order to reduce the unknowns and increase the predictability, well-targeted laboratory experiments, animal testing, and clinical trials must be run to develop the knowledge needed to put together a coherent picture of the disease process and the possible intervention(s). Each stage of this development process contributes to understanding the biological processes involved at the various levels of detail, from basic biochemistry, up through cellular biology, on into organs and systems, and finally to the symptom complex as it relates to patient perceptions. Experimentation and analysis works to build confidence in the proposed intervention and, hopefully, reduces some of its unpredictability.

Humans faced with the challenge of developing new drugs evolve mental models of the relevant biological

processes based on extensive experimentation and analysis. However, a disease process and the related biology are invariably complex, interrelated, and evolving over time. Even if all information about a process were available, it is still difficult for a human to effectively assimilate all of the data into a complete and coherent model or to identify all of its implications for disease progression. Thus, predictions based on such human mental models can be inaccurate and unreliable. Techniques that help to increase the knowledge and understanding of a particular disease, and that support the prediction of system behavior under specified conditions and over time, can reduce the uncertainty and risk involved with developing a new drug, and consequently reduce the time and cost.

One technique used to support the problem of predicting complex system behaviors has been the use of computer modeling and simulation. Numerous modeling and simulation methodologies exist. Most approaches used in the biomedical area have been either classic, numerically-based simulation techniques or knowledge-based, artificial intelligence techniques. These two approaches touch on the two ends of the knowledge spectrum needed to support drug development.

With numeric, or quantitative, approaches, sets of mathematical equations are developed that describe the behavior of a particular device or system. Such models are precise, but are only accurate if the device or system to be modeled is simple enough and understood well enough that it can be sufficiently and accurately described by a mathematical representation. Though such techniques have been used in numerous biological areas (Aliev and Saks, 1994) (Celada and Seiden, 1992) (Conolly and Kimbell, 1994) (Dalton et al., 1994) (Danis et al., 1995) (Hoang, 1995) (Pedley et al., 1994) (Robertson et al., 1994), the techniques do not scale up well to larger portions of human biology. They work well at the sub-cellular and chemical levels because the knowledge is usually available to model isolated portions at this level of biology with mathematical equations. However, such techniques are not well-suited for the higher, patient level models of the disease's symptom progression.

At the other end of the modeling spectrum are the heuristic-based approaches used to build knowledge-based and expert systems to model a particular disease at the patient level. Numerous systems have been developed over the years that can diagnose a particular class of disease given the signs and symptoms experienced by a particular patient along with any relevant test results (Shortliff, 1976) (Patil, Szolovits & Schwartz, 1981) (Weiss et al., 1978) (Chandrasekaran, 1983) (Miller, Pople, & Myers, 1982) (Kuipers, 1986). These systems have also been used to plan therapy regimens based on the diagnosis. Such approaches work well when the knowledge involved is at a higher, associative, heuristic, cause and effect level, but are not as effective as traditional techniques for modeling lower level, more detailed knowledge.

In the middle of the modeling spectrum are the qualitative modeling techniques developed primarily in the engineering, as opposed to the biological, domains (Bobrow, 1985) (Davis, 1983) (deKleer and Brown, 1984) (Forbus, 1984) (Genesereth, 1982) (Rieger & Grinberg, 1977) (Round, 1989). Many of these techniques have also been applied to modeling biology as a test of the approach (Uckun, 1992). These approaches provide a middle ground for representing a device or system explicitly as a set of behaviors when pure quantitative knowledge is not available or is not needed for the reasoning process being studied.

Thus, a large amount of modeling has been done in the biomedical field, but it has tended to touch primarily on only one point in the knowledge spectrum in any given modeling effort. Though some work has been done to link model-based and associative knowledge (Fink & Lusth, 1987), little work has been done to link the disparate levels of knowledge available at the sub-cellular level up through the cellular and system levels to the patient level in the biomedical arena (Karp & Frieland, 1989). In order to do this, a combination of modeling techniques must be used that can bridge the gap between quantitative models and heuristic models. The following sections discuss the various forms of knowledge involved in modeling the entire spectrum of biology needed to support drug discovery, describe how such a model is actually developed, and illustrate the technique with an example.

2 CHARACTERISTICS OF BIOLOGICAL KNOWLEDGE

Knowledge about disease processes is inherently complex, variable, and uncertain. As described in the initial section, it also spans numerous levels of detail, some of which are better understood than others. The various levels of knowledge needed to more fully understand how a particular chemical compound can influence a disease process are as follows:

1) Sub-cellular, chemical interactions that generate a local environment

2) Individual cell behaviors within an environment

3) Sets, or pools, of a specific cell type within an environment

4) Aggregations of different cell pools whose chemical products and behaviors create a particular environment

5) The environment generated by the aggregation of cell pools that manifests itself at the patient level as the disease's symptom complex

Notice that this list contains a feedback loop, since the environment listed in number one is aggregated to generate a global system environment in number four, which in turn influences numbers 2 and 3, and is manifested at the patient level in number five as the symptom complex. The complex interaction between the levels of biological knowledge makes it advantageous to represent and implement the knowledge independently within a hierarchy to accurately specify all relationships between levels.

To illustrate the types of knowledge included in this list, consider an example of otitis media. Otitis media is an infection of the middle ear. It is usually triggered by a viral infection of the epithelial cells in the nasal passages, as with a cold or the flu. As a result of the infection, the epithelial cell ciliary function is impaired and fluid builds up in the sinuses and associated passages including, potentially, the middle ear. This fluid environment is perfect for the proliferation of bacteria, which exacerbates the situation, eventually causing enough fluid to backup so that pain and swelling occur.

The knowledge about otitis media needed to build a disease model can be categorized according to the five level hierarchy described above. At the sub-cellular level (Level 1), the chemical environment consists of the classic pro-inflammatory cytokines such as IL-1, TNF, IL-2, IL-4, histamine, and numerous others generated by the pro-inflammatory set of cells. This chemical environment evolves over time based on the initial viral infection and inhibits the middle ear's drainage ability, which eventually causes the proliferation of bacteria. The cytokines and other chemicals found in the environment are the products of individual cells responding to specific chemical signals, a Level 2 phenomenon. Initially, the cells are responding to virally infected epithelial cells. Later, cells are responding to other cells that had initially responded in a cascade of events. Each of the individual cells can be aggregated into pools of the particular type of cell with a set of possible behavioral states. To model the cascade of events that represents the course of the disease, a pool of each cell type is necessary to represent the change of cells in the system from state to state (e.g. the probability, likelihood, or confidence that an epithelial cell is virally infected or virally immune). This represents knowledge at Level 3. The dynamics of the cell pools captures the progressive nature of the disease. Together, the pooled sets of cells produce certain chemicals based on cell states and the extracellular environment. This constitutes Level 4 knowledge. This, in turn, creates the extracellular environment, which feeds back to the various cell pools and generates the symptom complex, Level 5 knowledge.

The otitis media disease process just described is a superficial example of the kind of knowledge needed to properly and effectively model a disease in support of drug development. To truly model the process, sufficient detail must be provided so that the model can run as a simulation. For example, in order to model otitis media properly, sufficient knowledge of all of the cell types involved must be available, including what calls them into the area, what possible states they can be in, what environmental factors determine what state they are in, and what they do and produce when they are in any given state. Methods of acquiring and representing knowledge about diseases to build a viable model must surmount a number of obstacles and issues. These are as follows:

1) *The level of existing knowledge about a disease process and the basic supporting biology tends to be quite shallow*

   The current state of knowledge about a disease process often is relatively shallow, existing only as a set of abstracted rules that are empirically derived. Thus, researchers often do not have information about exact causal relationships and only have empirical data on associations. For example, researchers may know that a drug affects the symptoms of a disease but may not know how it directly affects the biology that causes the symptoms. They may know that a particular cell is involved in the disease process, but do not know all of its effects on other cells in the area. Thus, though representation of such a large amount of knowledge needed for a model of a particular disease can be problematic, obtaining the needed knowledge to make a model work can be even more difficult.

2) *The size and complexity of human anatomy and physiology is large and variable*

   The sheer complexity of the underlying biology of a disease process is staggering. Not only do

aspects of anatomy vary from person to person but the basic biological processes are greatly modified by the background and history of a particular patient. The magnitude of knowledge that must be acquired and encoded can be prohibitive unless efforts are made to simplify and aggregate wherever possible and appropriate.

3) *Biological systems are based on self-generated feedback and control and are not static, but evolving, adapting systems*

Biological systems work because of the self-regulating feedback. They are, fundamentally, control systems. As illustrated above by the levels of knowledge needed to model a disease process, feedback is the foundation of the process. Though current simulation techniques can handle feedback, the level of complexity and the pervasiveness of the feedback in biological systems is beyond what is commonly modeled in other application areas.

4) *What is known about a particular disease process is often uncertain and contradictory*

Although the best scientific minds may have participated in model development and the best, most current research available was obtained, there are generally still large areas of uncertainty in the knowledge. Experts must synthesize knowledge from a wide variety of sources to bring to bear on the current modeling effort. Those sources may only be tangentially related or may be unrepresentative studies that provide insight without confirmation. For example, research on a particular cell type's behavior under certain conditions may be used to support the model even though the data from the research is based on an entirely different disease process than the one being modeled.

5) *The knowledge on which the model is built is constantly changing and evolving*

The biological sciences are constantly uncovering new information about disease processes. One of the goals of a biological model is to synthesize that knowledge to express a comprehensive theory about the disease. Once the model has been built representing the theory, the theory can be tested, refined, and validated. As new information emerges, the model can be altered to incorporate it.

As a result of these issues, not only is the actual modeling process difficult in a biological area, but the knowledge acquisition process is highly complex, as well. Thus, although at the lower levels of knowledge certain mathematical techniques may be appropriate and sufficient, at the higher levels less restrictive methods of representation are needed that allow for the complex feedback interactions and the large uncertainties in the knowledge. Modeling techniques that are forgiving of the lack and imprecision of knowledge are needed to effectively develop useful models in support of drug discovery.

3 THE KNOWLEDGE-CENTERED MODEL DEVELOPMENT PROCESS

Model-based reasoning is the use of the structural, behavioral, and/or functional knowledge of a system in order to diagnose, predict, and explain its behavior over time (Uckun, 1992). Techniques have evolved out of the area of computer science known as artificial intelligence, specifically knowledge-based systems development and qualitative reasoning (Bobrow, 1985) (Davis, 1983) (deKleer and Brown, 1984) (Forbus, 1984) (Genesereth, 1982) (Rieger & Grinberg, 1977) (Round, 1989), and have been used to model areas either not amenable to traditional numerical methods or for which the level of accuracy of a numerical model was not needed. Knowledge-based systems development techniques are used to acquire knowledge about the underlying model of a domain and to convert this knowledge into an interactive computer simulation. The form of the knowledge in the model is qualitative and/or heuristic, representing entities and relationships by relative values rather than by precise amounts.

Formal definitions of quantitative and qualitative modeling exist (Cellier, 1991) (Clancy, 1989) (Fishwick, 1989a), as do definitions of knowledge-based systems (Waterman, 1986). Much discussion has occurred around what constitutes quantitative vs. qualitative modeling and why one would use qualitative modeling techniques (Cellier, 1991) (Fishwick, 1989b). Knowledge-centered modeling attempts to merge these various techniques into a single, multi-level, hierarchical model. The knowledge available for the model can be precise mathematical relationships between entities or can represent abstracted rules provided by experts. The level of knowledge can vary from a high degree of precision to a high degree of abstraction within a single knowledge-centered model. The reason for the use of qualitative or heuristic modeling techniques in the drug development area is pragmatic --- the knowledge is simply not available in the detail and completeness needed at all levels of the biology required to build a

working model. In addition, even if it were, the model would be so large and complex that aggregation at certain levels into qualitative models would still be required in order for the simulation to run in a reasonable period of time.

Knowledge-centered modeling has the following key attributes:

1) it is hierarchical,
2) it is modular,
3) it can model sufficient depth,
4) it can model sufficient breath, and
5) it is synergistic.

This technique allows for sufficient depth at the cellular and sub-cellular levels to model potential drug interventions at their source, while at the same time covering sufficient breadth of the relevant biology to model the symptom complex. Because of the hierarchical and modular nature of the resulting models, the simulations are sufficiently complex to allow for new and unique insights into biological processes.

The knowledge acquisition process used to build such a multi-level model results in the collection of many different pieces of data and information from many different sources. This process is time consuming and human intensive, often involving multiple domain experts over the course of many months. It follows both a top-down, structural approach and a bottom-up, data-driven approach. High level, general, textbook-oriented knowledge is collected as well as low level details on the latest in vitro laboratory experiments. The goal of knowledge acquisition is to determine which entities need to be represented in the model and the key relationships between the entities that drive the behavior of the system. Once the knowledge has been collected, the implementation process begins.

Implementation of a knowledge-centered model is hierarchical in nature. The process starts by identifying the symptoms that are of interest in the drug development effort. For example, in otitis media, the key symptoms are fluid behind the ear drum, pain, and fever. During the knowledge acquisition process, identified symptoms get mapped backwards into the biological environment that drives them and ultimately to the cellular and sub-cellular biology that causes the environment.

Once the territory is mapped with respect to the disease and its biological drivers, the top level of the model is segmented according to the major structural/functional components involved in the disease process. For example, in otitis media these major components might be the bacterial propagation and toxin productions, the middle ear area, the mucosal surface of the middle ear and eustachian tube, and the sub-mucosal connective tissue area.

Once the top level structural and functional relationships have been defined, more details are developed within each module to define its functionality. For example, within the mucosal surface of the middle ear module, a model of epithelial cells and their associated ciliary functions would be developed. This constitutes the Level 3, cell pool, knowledge in the model. It includes details on how virally infected epithelial cells behave, as well as what happens when they are exposed to pro-inflammatory chemicals from other cells responding to the viral infection. Finally, their behavior from exposure to bacterial outputs also has to be included. Such details must be worked through for every major component of the model. The knowledge is pulled not only from the set of human experts available to the project, but also from any literature that can be found on the subject, including any and all information on the specific cells and chemicals of interest. Again, this is done using both a top-down, structural and a bottom-up, data-driven approach.

As the knowledge is collected to the lowest level of detail needed, usually the cellular or sub-cellular level, implementation of the working model begins. Implementation is an iterative process in which levels of the model are built hierarchically, tested, and assembled into a working whole. Object-oriented techniques can be applied to speed the development effort and one hallmark of knowledge-centered modeling is how readily it adapts to this approach. The implementation process involves defining associations between entities and the qualitative relationships between them. These qualitative relations may be represented numerically and, at the lowest level of the model, usually are numerical values. However, these values are converted into fuzzy categories and it is these qualitative values that are actually passed around the model. Thus, knowledge-centered models include quantitative and qualitative values as well as heuristic associations.

Before a knowledge-centered model is complete, the basic biology resulting from the disease process that is now embodied in the model must be mapped back up to the symptom complex. Specific cell models are the bottom most layer in a hierarchical representation of the disease process at the cellular level. Anatomical features of the disease group the cells according to the point of their principle effects. Thus, some cells may have a large influence in one part of the anatomy but not in others. These groupings are necessary to localize effects

and establish appropriate feedback between interacting cell types. This organization and grouping occurs as a natural result of building a knowledge-centered model and directly supports the symptom complex mapping. For example, the production of a specific cytokine by the cells in a given region can be mapped to the level of a particular symptom of the disease, such as the effect of histamine on the level of tissue plasma and consequently the level of edema in the area.

Once the model is functional and has a symptom complex output, it can be run as a simulation, tested, and validated against available biologic reference patterns. These reference patterns define, based on available research data, what the patterns of cell behavior should be under the disease conditions. They also provide details on the course of the symptom complex for the disease process itself. The model behavior is compared at all possible levels to available reference patterns. The model is modified and tuned so that its outputs fit within reasonable variation of the biologic reference patterns. This provides a baseline for model behavior from which model predictions about unknown scenarios (e.g. novel drug interventions) can be extrapolated with some confidence.

A knowledge-centered model constitutes a coherent theory of the disease process and underlying biology being studied. Once the theory has been validated to the best available data, the model's behavior can be explored and assessed for a variety of purposes. In the area of identifying new potential intervention opportunities, various pieces of the biology represented in the model can be altered to simulate a particular drug intervention. For example, in the otitis media scenario, an inhibitor of specific mediators of inflammation can be represented by decreased amounts of these mediators, such as complement and prostaglandins. Then the resulting symptom complex as output from the simulation can be evaluated for change from the untreated state (e.g. fluid levels decrease). To better understand confusing or conflicting clinical trial data for a potential new therapy, the model input data can be set to mimic the clinical trial situation with respect to type of patient and form of treatment. The resulting scenario can then be simulated, the model behavior examined at various levels, and the effects on disease outcome evaluated. This information can then be used to support the design of more effective clinical trials.

4 AN EXAMPLE - A MODEL OF OTITIS MEDIA

The top level components of an otitis media model include a cell model of the bacteria that can infect the middle ear, models of cells at the eustachian tube surface, models of cells in the submucosal connective tissue, and a model of how the cells and cell products drive patient symptoms. Anatomically, these represent layers in the middle ear. The bacteria invade the middle ear, irritate epithelial cells at the eustachian tube surface, and the immune system responds by sending inflammatory and immune cells into the submucosal connective tissue layer to counter the infection. Some inflammatory cells also appear at the eustachian tube surface to directly kill bacteria.

The underlying biology at the cell pool, individual, and chemical levels (levels 3, 2, and 1, respectively) is developed within each of these major components of the otitis media disease process. Detailed models of each of the cells involved in the disease process are necessary to represent the effect of disease processes at the cellular level. In otitis media, the cell types modeled include mucosal epithelial cells, PMN's, macrophages, mast cells, basophils, eosinophils, and a number of other immune cells. These cell models each include knowledge about the relative number of a cell type, the potential cell states, and the chemical production rates of the cells in each of the potential states. Information about the signals in the system that move cells between states and/or cause cells to generate their products complete the representation of each cell type. So, for example, a mucosal epithelial cell is modeled as being in one of several potential states, including resting, virally infected, virally immune, proliferating, and activated. Each of these states represents a distinct behavior, such as a specific chemical production profile and responsiveness to other chemicals. For example, resting epithelials do not produce anything, but are susceptible to viral infection and subsequent change of state, resulting in viral, IL-1, and interferon production. Resting epithelial cells can also become exposed to interferon and change state to a virally immune cell that still produces nothing, but is no longer susceptible to the state change to virally infected.

In addition to the knowledge about the basic biology, knowledge is represented about how cell products drive patient symptoms to model the observable effects of the disease. For example, the state of a given epithelial cell defines the ciliary function for which it is capable. Taken as a whole, along with glandular secretions in the area, the mucociliary function of the entire population of epithelial cells can be calculated to determine their ability to clear fluid from the middle ear. This ability to clear, along with glandular production rates, determines whether or not fluid backs up into the middle ear, and at what rate, to cause the key symptoms of a middle ear infection.

Finally, knowledge about treatment effects at the cellular level is necessary to model possible interventions in the disease process. These effects must also be mapped through the biology up to the patient level in order to experiment with the effectiveness of various therapies. Thus, for example, an antibiotic will increase the death rate of the bacteria, reducing the bacterial load in the area and causing an alteration in the patient's immune response. This alteration in immune response is manifested in a modification to the numbers of cells in various states and, consequently, the level of certain chemicals. This then influences the symptom complex. Other potential therapies can also be modeled in a similar fashion, allowing exploration and assessment of hypothetical treatment scenarios.

5 CONCLUSIONS

Drug development involves first finding a key locus of intervention in the biology of the disease which, if affected, alters the course of disease progression, and then testing this intervention in the laboratory and in clinical trials to verify its effect and assess its safety. The pathology of a disease is often so complex that it takes years of research to discover a leverage point that provides a cure or at least relieves the symptoms. Current approaches to drug development concentrate on standard laboratory experimentation to generate hypotheses and animal tests to further evaluate the hypotheses. This is a very labor intensive and time consuming process in which a positive outcome is not assured. It relies on discovering an insight, which happens in due course rather than on a fixed schedule. Testing a proposed intervention once found also can take years of research to determine if its effect is significant enough and if its advantages outweigh any potential risks or side effects. The use of models to support this process can reduce time and risk by supporting exploration of the relevant biology in a consistent, coherent, and complete fashion. Such models can also serve as electronic laboratories for testing hypotheses quickly and with little risk. However, until recently, available modeling techniques were not powerful enough or flexible enough to capture sufficient portions of a disease process' biology to provide significant support. Knowledge-centered modeling provides this power.

REFERENCES

Aliev, M.K., and V.A. Saks. 1994. Mathematical modeling of intracellular transport processes and the creatine kinase systems: a probability approach. *Molecular and Cellular Biochemistry*. 133-134:333-346.

Bobrow, D.G., ed. 1985. *Qualitative Reasoning about Physical Systems*. MIT Press.

Celada, F. and P.E. Seiden. 1992. A computer model of cellular interactions in the immune system. *Immunology Today*. 13(2):56-62.

Cellier, F.E. 1991. Qualitative Modeling and Simulation: Promise or Illusion. *Proceedings of the 1991 Winter Simulation Conference*. WSC '91.

Chandrasekaran, B. 1983. On evaluating AI systems for medical diagnosis. *AI Magazine*. 4(2):34.

Clancy, W.J. 1989. Viewing Knowledge Bases as Qualitative Models. *IEEE Expert*. 4(2):9-23.

Conolly, R.B. and J.S. Kimbell. 1994. Computer simulation of cell growth governed by stochastic processes: application to clonal growth cancer models. *Toxicology and Applied Pharmacology*. 124(2):284-295.

Dalton, J.T., M.G. Weintjes, and J.L. Au. 1994. Effects of bladder resorption on pharmacokinetic data analysis. *Journal of Pharmacokinetics and Biopharmaceutics*. 22(3):183-205.

Danis, R.P., R.S. Moorthy, and J. Savage. 1995. Microhemodynamics of retinal collateral vessel formation. *Medical Hypotheses*. 44(2):103-109.

Davis, R. 1983. Reasoning from first principles in electronic troubleshooting. *International Journal of Man-Machine Studies*. 19(5):403-423.

deKleer, J. and J.S. Brown. 1984. A qualitative physics based on confluences. *Artificial Intelligence*. 24:7-83.

Fink, P.K. and J.C. Lusth. 1987. Expert systems and diagnostic expertise in the mechanical and electrical domains. *IEEE Transactions on Systems, Man, and Cybernetics*. 17(3):340-349.

Fishwick, P. 1989a. A Study of Terminology and Issues in Qualitative Simulation. *Simulation*. 52(3):5-9.

Fishwick, P. 1989b. Qualitative methodology in simulation model engineering. *Simulation*. 52(3):95-101.

Forbus, K. 1984. Qualitative Process Theory. Artificial Intelligence. 24(1-3):85-168.

Genesereth, M.R. 1982. Diagnosis Using Hierarchical Design Models. *Proceedings of the National Conference on Artificial Intelligence*. 278-283.

Hoang, K. 1995. Physiologically based pharmacokinetic models: mathematical fundamentals and simulation implementations. *Toxicology Letters*. 79(1-3):99-106.

Karp, P. and P. Friedland. 1989. Coordinating the Use of Qualitative and Quantitative Knowledge in Declarative Device Modeling. In *Artificial Intelligence, Simulation, and Modeling*. L.E. Widman et al. (eds.). John Wiley & Sons.

Kuipers, B. 1986. Qualitative Simulation. *Artificial Intelligence*. 29:289.

Miller, R.A., H.E. Pople, and J.D. Myers. 1982. INTERNIST-1: An Experimental Computer-based Diagnostic Consultant for General Internal Medicine. *The New England Journal of Medicine*. 307:468-476.

Patil, R.S., P. Szolovits, and W.B. Schwartz. 1981. Causal understanding of patient illness in medical diagnosis. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*. 893-899.

Pedley, T.J., et al. 1994. Gas flow and mixing in the airways. Critical Care Medicine. 22 (9 Suppl):S24-36.

Rieger, C. and M. Grinberg. 1977. The Declarative Representation and Procedural Simulation of Causality in Physical Mechanisms. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. 250-256.

Robertson, M.J., et al. 1994. Agonist-antagonist interactions at angiotensin receptors: application of a two-state receptor model. *Trends in Pharmacological Sciences*. 15(10):364-369.

Round, A. 1989. Knowledge-Based Simulation. *The Handbook of Artificial Intelligence*, Vol IV. Barr, Cohan, and Feigenbaum (eds.). Addison-Wesley.

Shortliff, E.H. 1976. *Computer-Based Medical Consultation: MYCIN*. Elsevier, New York.

Uckun, S. 1992. Model-Based Reasoning in Biomedicine. *Critical Reviews in Biomedical Engineering*. 19(4):261-292.

Waterman, D.A. 1986. *A Guide to Expert Systems*. Addison-Wesley.

Weiss, S.M., C. Kulikowski, S. Amarel, and A. Safir. 1978. A Model-Based Method for Computer-Aided Medical Decision Making. *Artificial Intelligence*. 11(2):145-172.

AUTHOR BIOGRAPHIES

PAMELA K. FINK is Executive Director of Discovery and Technology Systems for Medical Science Systems, Inc. She received a B.A. degree in Mathematics and Philosophy from Eckerd College in 1989 and a Masters' and Ph.D. in Computer Science from Duke University in 1991 and 1993, respectively. Her research interests include issues in knowledge acquisition, representation, and utilization as they relate to problem solving and computer software implementation. Currently, she develops computer models of complex biological processes to support pharmaceutical and medical device research and development. She is a member of the American Association for Artificial Intelligence, the IEEE Computer Society, the Association for Computing Machinery, and the Society for Computer Simulation.

L. TANDY HERREN is a Staff Scientist for Medical Science Systems, Inc. She received a B.S. degree, an M.A. degree, and a Ph.D. in Psychology from Ohio State University. She also received an M.S. in computer science from Ohio State. Her current research interests involve biomedical simulation for therapy development, and knowledge acquisition and representation methodologies. She is a member of the American Association of Artificial Intelligence and the IEEE Computer Society.