Using Heuristic Models to Bridge the Gap Between Analytic and Experimental Models in Biology

Glen E. Ropella[1,2]**, C. Anthony Hunt**[2] **and Dev A. Nag**[1,2]
[1]**Tempus Dictum, Inc.,** [2]**The BioSystems Group, University of California**
513 Parnassus Ave., San Francisco, CA 94143-0446
gepr@itsa.ucsf.edu   hunt@itsa.ucsf.edu   dev@tempusdictum.com

Abstract

Models are often classified based on their structure as opposed to the context in which they are developed or found useful. However, a model's development is driven by the context in which it is expected to be used. Because most computer models are developed with low level programming languages, those models often end up being over fitted to a single, unique, context and fail to be useful in other situations. So doing severely limits the model's usefulness. This problem exists in all domains where models are used; but, it is particularly prevalent in the modeling of biological systems because biological systems are always changing and their behaviors can be more sensitive to context than models in other domains. Multiple, separate models of a biological system are required to begin adequately representing that system's behavior. This is in stark contrast to the biological components, themselves. They function in many different situations, whereas in many of those situations their corresponding models would fail. This problem presents a fundamental breakdown in the extent to which any computer model can represent its biological referent because there is, simultaneously, a sensitivity to context and a robustness to changes in context that exists in the referents but not in the models. In this paper, we present the basics of a modeling method, FURM (Functional Unit Representation Method), that attempts to address this breakdown by enforcing and encouraging the use of some basic methodological model development principles. We briefly present a constructive, and therefore heuristic, model of an isolated, perfused rat liver as a first article demonstration of the method.

1. INTRODUCTION

What is meant by the word "model"? The etymology generally indicates that "model" is closely related to "measure" [American Heritage 2000]. Taken in that sense, a model can be thought of as a device against which to measure an artifact or phenomenon. This definition consolidates the diverse uses of the word from fashion models to geographical maps to systems of mathematical equations.

With this definition in mind, what is it we do when we model some process, artifact, or system? What are we trying to create? The biggest element of this question lies in how the model will be compared to the referent. What aspects of the referent are observable? What makes those observables salient? Which of the observables lend themselves to the ascription of quantity and which are more qualitative? This line of questioning leads us to the first consideration in any modeling effort: use cases. [Jacobson et. al. 1999, Wiley 2000, and Beck 2000] The first consideration in modeling is to determine why the model is being created and the situations in which the model will be used. The use cases tell us what the model is good for, what it must do, who it must serve, how long it must be useful, etc.

Use cases elicit the various aspects [Kiczales et. al. 1997] or facets [Zeigler 1984] of the newly posited device. In order to do this, a common sense approach would be to elicit the salient usage situations of the referent. Often, however, it is very difficult to delineate the salient aspects of a biological object because they can appear to have a (dense) continuum of uses and situations in which they are relevant.   In fact, it is a well-known theorem in metamathematics called Tarski's Theorem that "assuming that the class of all provable sentences of the metatheory is consistent, it is impossible to construct an adequate definition of truth in the sense of convention T on the basis of the metatheory"[*] [Tarski 1956], which leads von Neumann to claim that for particular types of systems, their description will be indefinitely long. [von Neumann 1966] This forces modelers to choose, somewhat arbitrarily, the aspects of their model from the potentially very large set of possible properties or attributes the model might have. One consequence is that the modeler is now open to the rhetorical fallacy of "begging the question" (*petitio principii*) which in modeling becomes the fallacy of "inscription error" [Smith 1996].

Given that any modeling effort may be faced with a huge space of relevant use cases, how does a modeler choose which use cases are salient? The answer depends on

---

[*] Convention T simply states how statements in the metalanguage map to statements in the language. It is interesting to note that Tarski, unlike Gödel, treats colloquial language and the philosophical understanding of truth as well as formal systems or computer programs.

the core context in which the model is expected to be used. This core context provides the impetus for the development of the model, sets the expected lifetime of the model, defines the actors involved in its development, usage, and maintenance, and is the source for all the model's requirements. The core context will usually be ambiguous enough to preserve model extensibility while providing the basis of a guideline for development decisions.

The core context can be as broad and ill-defined as an entire domain like "physics" or as particular as a single very specific process like that of the canonical traveling salesman. The key property of the core context lies in the modeler's ability to collapse it into a detailed enough prescription for what the model must do. This reduction of the context is particular to the modeler (or team of modelers) and implies that the model is developed through subjective interpretation of the modeler's perception of the context.

Hence, modeling is a subjective enterprise. It has always been and will always be subjective. And when viewed this way, it is easy to see that modeling efforts are epistemological and cognitive processes. They are not ontological. Models tell us absolutely nothing about the world, except in what they tell us about the psychology, reasoning, and sociology of the humans engaged in the process. Even in the extreme cases where a model is fully accepted by a very large population of people, the models remain, inherently, cognitive. For example, up until 1960, the concept of the distance of a meter was modeled by two marks on a metal rod kept in Paris, France. Everyone in the world (who had reason to refer to things like "meters") accepted that this metal rod was an adequate model of the meter, due, no doubt, to many meetings and much argumentation about the composition of the rod, what temperature to keep it at, etc. This type of model becomes accepted and completely subsumed by the knowledge and behaviors of the people who accept it so that discussions of it are relegated to obscure subcultures, anthropologists, and standards organizations. Although models like this are treated as if they are true and represent some form of ontological knowledge, they remain models and they are nonsensical when separated from the referent they are intended to model.

There are many distinct ways in which modeling is used in the study of biological systems. A taxonomy of biological models can be organized based on the various contexts in which the models are used. Figure 1 is an illustration of relative relationships between various categories of models, methods, and devices used in biomedical research, circa 2000. The focus here is on organisms. Most animal research (biological analysis) is done with model organisms (MO). Frequently, they are specialized inbred strains. Patients and animals not bred for research are outside this diagram. *In vivo* experiments are done on intact, living animals. There are two categories of *in vitro* biological models (BMs). The first uses living parts taken from model organisms to build (synthesize) artificial systems that meet some objective. Examples include *in vitro* cell, tissue and organ cultures. The second comprises
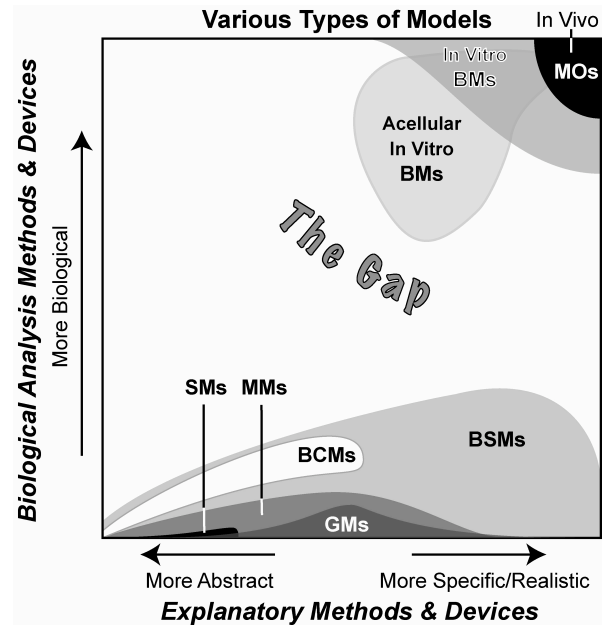
Figure 1. Model types arranged according to abstraction level versus biological character. The more abstract models indicate a higher capability for simple and focused representation. The more realistic models indicate a higher capability for aggregating collections of facts. The biological axis indicates the degree to which a model resembles, in detail, its biological referent. MO: Model Organisms (supplies for in vitro experimentation); BM: Biological Models (in vitro cell, tissue, and organ cultures and standardized cell forms); BCM: Biochemical Models (studying abstracted, specific biochemical events and processes); StM: Statistical Models (explaining or accounting for data taken from other models); MM: Mathematical Models (biomimetic functional models used for prediction); GM: Graphical Models (heuristic descriptions used for clarity and communication); BSM: Biological Systems Models (partially heuristic, partially predictive, biomimetic functional and structural models used for evaluating canalizing aspects of the referent); NCMs: New Class of Models (heuristic, biomimetic models used for evaluating explicit hypotheses in the context of many aspects of the referent).

standardized single cell forms, such as E. coli, the yeast S. cerevisiae, and newer, "synthetic" organisms [Ferber 2004]. Figure 1 focuses on the first category but does not exclude the second. *In vitro* BMs are created, constructive systems. Specialized devices, support instruments, and reagents are required to keep the living components stable and functioning. As a consequence of their *in vitro* nature, new properties (biological, mechanical, chemical, etc.) are imposed onto all *in vitro* BMs and some original properties are lost or altered. Relative to referent model organisms they are less realistic and less biological.

Acellular *in vitro* BMs are constructed starting with living parts, typically isolated cells. They are broken (killed) and components thereof are taken for study *in vitro*. Examples include cytoplasmic fractions, organelles, nucleoplasm, reconstituted metabolic networks, etc. Many of the original biological features are absent. Linkages between networks and modules are severed. Nevertheless,

the retained features are thought to mimic many corresponding features in the original organism. Biochemical models (BCMs) are assembled from individual synthetic or purified biological components to study specific events and processes. Relative to acellular *in vitro* BMs these synthetic BCMs are simple, abstract systems. They contain little if any biology. Typically, all the variables and their values are known.

Of course, the several classes of inductive analytical models that are used in biomedical research have no organic components. Statistical Models (StMs) are not actually models of biology, but they are ubiquitous in biomedical research, and are referred to as models. They focus on data and much less so on the biological processes and events that give rise to the data. They are models that are intended to explain or account for the variance structure in the data. Mathematical Models (MMs) in biomedical research encompass a huge literature. They are all orders of magnitude simpler than *in vitro* acellular models and even further removed from *in vitro* and *in vivo* models. Typically, inductive analytical MMs represent abstract, very useful metaphors of specific processes and events that occur *in vitro* or *in vivo*. They are predictive theories, and as such they too focus on the collected data. Their measure of success is how often and how well they fit the data. Graphical Models (GMs) include formal graphs and the informal diagrams that are intended to capture features of a researcher's mental model of some aspect of biology, such as her or his vision of the factors and influences involved in transport of a solute into and out of a cell. Examples of GMs that fall in between these extremes include cognitive models of metabolic and gene regulatory networks.

Biological Systems Models (BSMs) are cognitive models grounded in a reductionist approach to biology. Systems of equations (typically differential equations) and/or probabilistic networks, to name just two approaches, are used to represent or explain known and hypothesized phenomena. As with MMs, BSMs begin with a long list of simplifying assumptions that dramatically distance them from the *in vitro* and *in vivo* BMs that were the sources of the modeled data. To understand biology we need to systematically reduce the list of simplifying assumptions with the long-range goal of striving to eliminate it. The inductive analytical models are simply not adequate hosts for the specific new types of data being acquired in the many domains of biomedical research. We need to understand how the parts function within a biological whole. This opens the door to new synthetic models and methods, including agent-based models, for understanding biological systems. Realization of this opportunity will represent a new class of models that can fill *The Gap* in Fig. 1, as illustrated by the dark gray area in Figure 2. We need these new approaches to bridge the gap in understanding between *in vitro* BMs and current MMs and BSMs. These new approaches will need to work, primarily, by decentralizing the modeling process (into modules and components that can be plugged together in different ways)
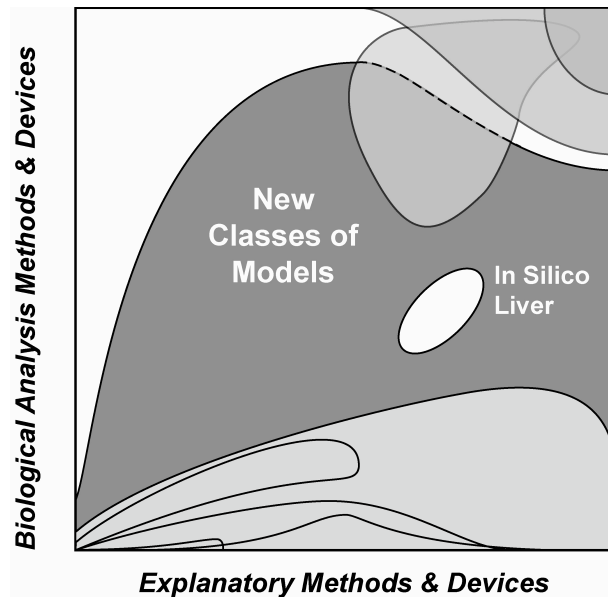
Figure 2. An extension of Figure 1 showing a New Class of Models (NCMs) that fills *The Gap*. The in silico liver is a first article example for this type of model and fits squarely within *The Gap* but does not yet provide a full bridge.

without requiring that all the referent data be of a specific type or that it even exist.

The new class of synthetic simulation models (NCMs) identified in Fig. 2 will include agent-based models. As these NCMs become validated, they will begin to provide a structural and interactional framework for assimilating genomic and proteomic data. Hybrid System Models are needed to systematically and iteratively test and revise prevailing ideas about how biological systems and subsystems actually function together. The expectation is that they will do so by virtually reassembling virtual parts and components into a networked system of systems that can account for an increasing fraction of the accumulating data in combination with a shrinking list of assumptions. Clearly, some of these new NCMs will have more in common with *in vitro* and acellular biological models than with their predecessor MMs and BSMs. That is the whole idea. Their design and construction are expected to draw extensively from continuing advances in computer science and complex systems science. Whereas a goal of earlier MMs was to offer fits to data and plausible predictions, the goal of the new class NCMs will be to function as exploratory experimental systems (synthetic in silico constructs) that can be used iteratively to narrow the list of competing explanations about how biological systems function. One example is the In Silico (rat) Liver shown in Figure 2; it is briefly discussed at the end of this paper. It and related models currently under development within the UCSF BioSystems Group are some of the first realizations of the new class of models.

2. INDUCTION VS. SYNTHESIS

The categories of models in Fig. 1 are distinct but not disjoint, and the usage contexts are not mutually exclusive. In fact, they are vague and open to interpretation by the modeler. The purpose of laying them out is intended to help a biological modeler make decisions about what methods and technologies to use to achieve particular ends. As discussed before, those ends are best achieved by laying out the salient use cases, which can now be organized and determined in relation to the various contexts and model types discussed above.

The core context of a model guides the development of any tools and techniques the modeler might find useful for constructing models. This typically is an ad-hoc collection of techniques that the modeler has previously been exposed to and has found useful. And some disciplines engage in explicit attempts to provide a decision-making structure to help choose the appropriate methods, e.g. domain analysis in systems engineering. The methods used by modelers vary a great deal from very precise methods like those recommended by the Unified Modeling Language (UML) to very ad-hoc tools like transient drawing on white boards. All these methods are relevant and often critical elements to creating useful models.

When building a model to adhere to multiple use cases, some of those cases can interact in a way that is either contradictory or difficult to reconcile. In many situations there are methods to deal with conflicts, e.g. [Walker 2003] and [Lin 1994]; but, sometimes the conflict is fundamental and has to be resolved simply through the expertise of the modeler. One such fundamental conflict is that of the difference between the inductive and synthetic methods for building systems.

Today, the biological models in the lower half of Fig. 1 are usually built by taking data from an experimental system, analyzing that data, creating a mapping between the structure of the experimental system and components of the data, then representing those components of the data with mathematical equations. Those equations are then implemented, executed, and validated against the data. This method relies heavily on psychological induction and in some cases on the types of induction achievable by pattern recognition algorithms.

The inductive method [Steels and Brooks 1995] for creating models stays very close to the data and, when successful, provides models that extrapolate beyond the original data from which the model was induced. This makes it ideal for the development of predictive models.

Although this umbrella is very broad, it is instructive to contrast this method with constructivist or "synthetic" methods. The synthetic method is common in non-biological domains and appears to have been used some in biology, as well, particularly when constructing mechanical or electrical analogs of biological units and the *in vitro* models in the upper right of Fig. 1 (and some *in vivo* models, as well). The synthetic method consists of proposing and constructing building blocks that can be assembled to create an artificial system that functions in the real world. This artificial system or analog is exercised and measured in the same way as its referent system. Then the data taken from the analog is validated against the data taken from the referent. The essential difference between the two methods lies in the fact that the inductive method explicitly uses the observed phenomena (the data) as its input whereas the synthetic method starts with proposed building blocks and the relations between building blocks. A useful way to think about this distinction is that the inductive method takes the ontology defined implicitly in the data as an inherent assumption and the synthetic method attempts to construct an ontology (and an analog that "lives" in that ontology) that can realize the data.

Another useful way of thinking of the distinction between the two methods is to think of a mapping from the space of generators or mechanisms to the space of phenomena. The inductive method starts with the phenomena and works backward to the generators in an attempt to discover the inverse mapping from range to domain. The synthetic method, in contrast, works forward from domain to range. When phrased this way, it becomes clear that neither method is superior and both should be used when modeling. Figure 3 provides a visual depiction of how these two methods relate.

Inductively generated models have two primary benefits over synthetically defined models. The first, mentioned above, is their consistency with data taken from the referent. The tolerance of this consistency can be measured through similarity measures [Ropella et. al. 2003] acting on the referent and model data. This consistency, with a tight enough tolerance, can provide a somewhat trustworthy extrapolation into regions of behavior space that are not covered in the available data, i.e. they are predictive. The second benefit is just as important and speaks to the feasibility of any modeling effort and is practically more important, in some ways, than prediction. These models compress the complicated dynamics and structure of the real system into a much smaller and clearer representation that (typically, equations or systems thereof) allows us to simulate reality faster than real-time and with relatively few resources.

However, the primary limitation of inductive modeling is its inherent dependence on the necessary epiphanies, insight, and experience of the modeler and the dynamics of the people analyzing the data, publishing their experiments, results, and conclusions, and proposing the salient components of the model. There is no brute force (algorithmic) way of inferring the structure of the generator system from the observed data. Inductively developed models are lossy compressions in the same way that the observed data is a lossy representation of the system.

The primary benefit of the synthetic method is a result of the difficult and iterative process of actually designing and building an analog that behaves similarly to the referent. The tacit experience the modeler gains while going through this process and the potential behaviors of the analog in its

4

finished state yield not only the potential for extrapolation, but provide the modeler with a detailed control surface (parameterization) that allows the analog to explore new, possibly fundamentally different regions of behavior space. These different regions may or may not be part of the referent's behavior space; but this exploration can be helpful in understanding the structure and dynamics of the referent.

Of course, the fundamental difficulty with the synthetic method is establishing requirements for building an analog that functions at all like the referent. This requires knowledge of the function of the referent, of plausible mechanisms for that function, and of relevant observables by which the analog and the referent will be measured.

Contrasting these two methods, which every modeler does to some extent, is useful for determining how a model can best be put to use. With this in mind, we return to the original questions: What are we doing when we engage in the process of modeling? What are models used for? How do we purposefully choose a modeling method that achieves our ends?

3. HEURISTIC VS. PREDICTIVE MODELS

All of the above points (intuition, compression, salient aspects, aspect resolution, model type, etc.) make it obvious that modeling is more art and engineering than it is a science. Building models that satisfy some set of requirements requires many levels of expertise and interdisciplinary interaction. Moreover, there is no curriculum or textbook that a neophyte modeler can study in order to become a good modeler. Good modelers become so as a consequence of training and experience. There are no shortcuts to that practical skill. If one wants to be a good modeler, one simply has to build lots and lots of models.

However, it is our belief that the synthetic or constructive method significantly contributes to a modeler's understanding of what it means to model and how best to go about creating effective models. One of the best ways to learn about a phenomenon is to try to build a device that exhibits the behavior your looking for. This principle is used in the design of everything from children's toys to electric circuits. And it seems to have resulted in a beneficial side effect of von Neumann's constructive efforts to build reproductive machines. [McMullin 2000].

With this in mind, it is a small jump to see why many biological models are heuristically impoverished. Most biological models in the lower half of Fig. 1 are inductively defined mathematical compressions of the phenomena they are intended to model rather than being constructive models used as deep analogs of their referents.

There are many aspects of biological systems that cannot now be adequately modeled. The primary specific failure of biological models we expect to address through constructive modeling is robustness. Biological models are fragile to the context that drives their development. For example, a specific ODE model formulated to replicate the data taken from a specific set of biological experiments is rarely applicable to another set of experiments without
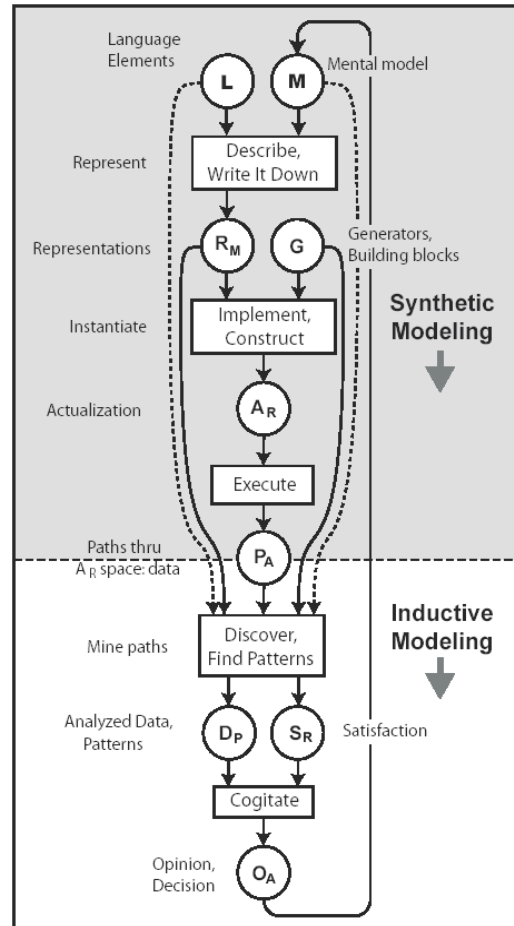
Figure 3. Conceptual model of the synthetic and inductive modeling processes showing that they are, at heart, the same process but with different initial and final states.

reformulation. That reformulation may be minor (adjustment of parameter values) or major (addition of new terms or expansion to a system of equations); but it is reformulation. The old model is considered inadequate because it is a lossy compression of the data from the first set of experiments. It did not capture the generative mechanism of the biological system in order to be reused to account for the new experimental data. All of these models are "broken" in some deep sense. They obviously do not model their referents very well.

The fact that current biological models are this impoverished is well known but often overlooked. The point was recently raised by Dr. Sydney Brenner at the NIH BISTI 2003 Symposium: "The man in the street doesn't believe in evolution, because he says we're trying to tell him you can take a black and white television set, you can make random mutations in it that will turn it into a colored television set. He knows if you tamper with a television set, the most likely thing is you'll break it. So how do biological organisms not get broken all the time? That tells

us there must be an architecture within the way they are constructed logically and actually which makes them resistant to this and allows evolution to proceed. And that will be a feature of what we're looking at. It will be there; it will be underlying everything that we do."

Discovering and understanding this robust architecture is the goal of heuristic biological models. And the synthetic method is one technique for building such models. In that same keynote talk, Dr. Brenner also makes the connection between the synthetic method and discovering this robust architecture by pointing to von Neumann's constructive proof of self-reproducing automata. Constructive proofs are inherently heuristic.

Of course, the problem with constructive proof and, by extension, the synthetic method as a whole, is coming up with devices that successfully generate the behavior sought. Take that consideration along with the prior point that good modelers emerge out of training and experience and it is easy to realize that what we need are methods for facilitating the building of models, rather than methods that help build better models. The focus will not be on the resulting models. The focus will be on the process of model building. The more models we can build and evaluate, the more likely we are to find plausible mechanisms underlying the robust architecture of biological systems.

4. THE CASE FOR A NEW METHOD

We have attempted to lay the foundations for the argument for a new modeling method. Carrying forward from the propositions that:

- effective models are developed to stay close to their context;
- models are purely subjective, psychological artifacts;
- the way to be good at modeling is to build many models; and
- heuristic models are constructive analogs of their referents;

we can tentatively infer that modeling efforts (especially of complex biological referents) should be "thought experiments" where many different models are constructed out of many sets of primitive building blocks and evaluated against some criteria defined by the context.

We are engaged in an attempt to establish whether this conclusion is sound or not by building a modeling method, tools that support the method, and clinically relevant models that adhere to the method.

To place this effort into the model taxonomy above, we are further proposing that these NCMs can fill a gap in the way biological models are used in the study of biological systems. The NCMs are developed using the synthetic method to explore the space of possible mechanisms, biotic and abiotic alike, that can generate the salient aspects of biological systems. The gap is one of **exploratory modeling** in the tradition of von Neumann's automata [von Neumann 1966], Ulam's description of "experiments in theory" [Ulam 1960], and monte carlo methods [Metropolis

and Ulam 1949] and which extend to more modern efforts like the "opaque thought experiments" of [Di Paolo et. al. 2000], Mark Tilden's BEAM robots, Danny Hillis' Tinker Toy computer [Dewdney 1989] , and Gerald Edelman's Darwin automaton [Krichmar and Edelman 2002]. These efforts are united in the emphasis they place on the synthesis of and experimentation with many various models.

5. FURM

The modeling method we have developed is one compiled out of behaviors learned to be effective through our modeling experience. The guidelines that compose the method are subject to change and are somewhat ambiguous as befits methodology.

Because of the overwhelming tendency of biologists to see biological systems, at the selected level of resolution, as being composed of functional units, we have chosen to center the method around the development of functional unit analogs. This means that the models that are developed will adopt a perspective that biological units do exist and have some form of unity. Hence we call the method the Functional Unit Representation Method (FURM).

This underlying assumption of unity allows us to effectively use Object-Oriented (OO) technologies as the basis for the constructive models. However, OO does not generalize well into biology because functional units do not have hard boundaries and often exhibit autotelic properties. So, we violate some of the strict OO guidelines and adopt Agent-Based Modeling (ABM) techniques where appropriate. We suggest four fundamental guidelines:

- standardize interfaces to multi-paradigm, multi-mode; and trans-domain (a.k.a. cross-trophic) models;
- use discrete interactions;
- design for an extended life cycle; and
- define observables that will submit to a similarity measure.

For our modeling functional units, the preceding four guidelines flow down into nine more concrete behaviors (principles).

1. Iterative Modeling - Because of the nature of biological systems, there can be no optimal or finished model. Design and build so that models can undergo continual evolution.
2. Model Comparison/Contrast - Isolated models are nonsense. They only make sense in relation, comparison, and contrast to other models.
3. V&V - Every component of a model should be open to validation (data permitting) based on clearly defined measures. Verification should be limited to the examination of networks of validatable components.
4. Aspect-Oriented Modeling - Models must be capable of representing multiple, possibly incommensurate, perspectives.
5. Arbitrary Functional Granularity - The composition of any biological system can be dynamic, and so it is open to interpretation. Hence fixed compositional attributes

(like a specific networked hierarchy) will always be weak points in a given model. If a model preserves the ability to re-specify or re-interpret its functional granularity, these weak points are hardened.

6. Experimental Procedure Encapsulation - Capture experimental procedures in an unambiguous way so that different experiments give similarly formatted results and are easily repeatable, preferably automatically repeatable, and extensible.

7. Multiple Models - Prefer running multiple models of the same referent in tandem.

8. Automated Model Generation - Build mechanisms to automate (or at least assist in) the generation of new models.

9. Experimental Computing - Treat any computational process as if it were a black box with an impenetrable boundary around it.

FURM is intended to do two primary things:

- Encourage disciplined experimentation with digital computers, including the, still rare, discipline of keeping track of source code as a type of experimental material like a chemical compound or a device, and

- Bring enough rigor to the modeling process to allow us to automate the generation and evaluation of models.

If we make progress on either of these goals, we will consider that success as supporting evidence for the above inference that modeling should be a thinking tool. The method, as it stands, is enough to make progress on the first goal. If it is adopted and used regularly by even a small group of biological modelers, then it will push the community further in the direction of disciplined exploratory modeling. Evidence of success will depend, fundamentally, on whether the models we generate through FURM are successful in their respective domains. For example, the ISL described below presents us with a facility for evaluating heuristic or predictive hypotheses about pharmacokinetics.

The second goal, however, requires an additional component that has not yet been developed. One or more generative mechanisms must be created from which to automatically generate new models. Our current plan is to use the models we develop with knowledge acquired directly from experts and domain literature to infer a collection of building blocks and a syntax for plugging those building blocks together to create models. Then we intend to apply an evolutionary algorithm, with our similarity measures providing the selection mechanism, to the models so generated. The successful models can then be analyzed to provide researchers with brain food to help them think about and design experiments on the referent systems.

A common problem in evolutionary computing is the generation of unintelligible (and non-analytic) individuals. Because the building blocks and syntax will not be a general purpose computing language and because the method centers on functional units, we do not expect to be faced with non-analytic models.

In the end, we expect the system to be a biological sandbox in which a biologist (or modeler) can discover innovative explanations and propositions for how the real biology might give rise to the phenomena embodied in the similarity measures. Ideally, it would operate as a biologist's assistant.

Brief mention also needs to be made of other "soft computing" techniques that may help us to achieve the above listed goals. There seem to be many places where AI techniques like fuzzy systems, ANNs, Bayesian networks, etc. can augment the basic method in order to make automatic model generation and selection steadily more efficient. It is our hope that these techniques can both constrain and expand the system in order to help the modeler balance the healthy tug-of-war between exploration and exploitation in science.

6. THE IN SILICO RAT LIVER

A liver is built of several lobes, each made up of hundreds of lobules. Each lobule consists of several acini organized around terminal afferent vessels. See Figure 4. An acinus is a network of sinusoids, tubular structures that carry blood from the portal vein (PV) to the central hepatic vein (CV). The sinusoids separate mesh works of one-cell thick plates made up of hepatocytes. Intra-acinar hepatocytes exhibit properties that are location-specific, often identified by different zones. To test and challenge our current concepts about liver function and the role of hepatic microenvironments in normal and disease states, we need models that can represent a liver as an organized assembly of individually distinct primary units, either acini or lobules. The new models need to be sufficiently flexible
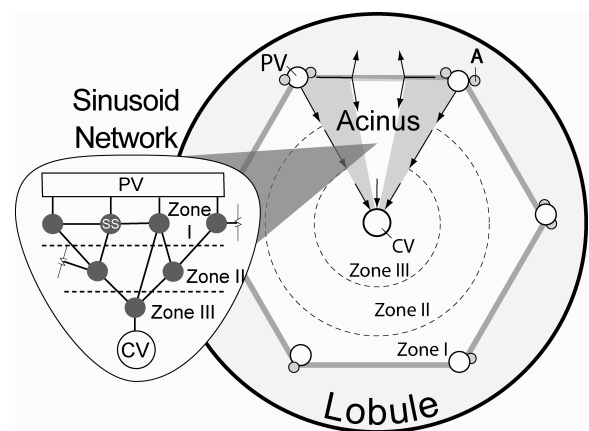
Figure 4. An illustration of general hepatic lobule structure. Lobules are comprised of several acini, a network of sinusoids connecting the terminal portal vein tract (PV) to the central hepatic vein (CV). Within the model the sinusoid is represented by a directed graph with nodes in three zones and agent structures called Sinusoidal Segments (SS) placed at each node. A: terminal arteriole.

to represent different aspects of hepatic biology at several levels of resolution. Our in silico liver (ISL) is such a model.

The *in situ* liver perfusion protocol that is the source of our experimental data is detailed in [Roberts and Anissimov 1999]. The experimental data are outflow profiles of fractions of an administered compound. Data for several compounds are available. The liver interacts differently with each compound. In some cases the compounds are extensively metabolized.

Two in silico system models are implemented. RefModel is the accepted, reference mathematical model (for details see [Hung et. al. 2001]). ArtModel is our functional unit model. Isolated parameter vectors are chosen based on some criteria for likely solutions. Input parameters and initial conditions are iterated until some stopping criteria are met. The models are run with those parameter vectors. Output is logged over time. Measurements are taken and plotted in conjunction with the experimental *in situ* data. The plots from ArtModel are identified as being close, or not, to the corresponding *in situ* data using a similarity measure.

The explicit hypothesis being tested by the in silico experiments is that the selected parameter vectors cause the model to generate output that is experimentally indistinguishable from the *in situ* data. The similarity measure can be used to automate the evaluation of the solution sets put forth by the models. Those results make possible automatic searches of the parameter space for regions that, together with the events elicited by the models, solve the problem (e.g., match the *in situ* data).

The assumptions made, especially those represented by the experimental data, are too numerous to delineate. So we focus on those deemed most salient. At the highest level we assume that rat livers are roughly equivalent and that the perfusion technique and experiments do not push the liver into pathological states. The RefModel assumes that both liver vascular components and the contributions of the various pieces of the experimental apparatus can be pooled (see [Roberts and Anissimov 1999] for details). The ArtModel assumes that liver function, as a whole, is an aggregate of acinus function, that sinusoids are primarily vascular objects, and that transit time for perfusate is governed by stochastic interactions between various agents inside the vascular structures in combination with a perfusion pressure at the inlet catheter and PV. At a more detailed level the four primary assumptions are:

- Outflow profiles alone are lossy projections of liver behavior. Models induced from such profiles are not rich enough for knowledge discovery. More accurate physiological models are necessary to begin fully exploring the liver behavior space.
- Hepatic vascular structure can be represented by a directed graph.
- The primary functional unit is the acinus or composites thereof.
- Outflow for an inert, polar solute such as sucrose is solely a function of the extracellular (vascular cavity) space, its geometry, and properties.

At present, we use a simple interval similarity measure. A set of experimental outflow profiles is used as training data. From this data, we calculate a distance, D, from a reference that will be the basis for a match. We then take two outflow profiles and pick one to be the reference profile, P^r. For each observation in P^r create a lower, P^l, and an upper, P^u, bound by multiplying that observation by $(1 - D)$ and $(1 + D)$, respectively. The two curves P^l and P^u are the lower and upper bounds of a band around P^r. The two outflow profiles are deemed similar if the second profile, P, stays within the band. The distance D is the standard deviation of the array of relative differences between each observation and the mean observations at that time. For the training set we choose repeat experimental data on the same subject to calculate D. This allows us to arrive at an estimate of intrasubject variability.

Results from the simulation along with additional details can be seen in [Ropella and Hunt 2003].

7. DISCUSSION

FURM is intended to assist the computational biologist in bringing more discipline and more of the accumulating, research-derived information to the upstream development of widely different models and to assist in exploring the range of behaviors of any given model and its referent system.

Scientific modelers often restrict themselves in the methods they use and in the data on which they focus. Their models often assume many simplifying properties (like linearity) to make them easier to use, understand, or teach. At times, sheer psychological inertia prevents new or counter-intuitive models from gaining ground or being perceived as credible. When methods are firmly based on validation (or invalidation), their origin does not matter. What matters is whether or not they are useful.

FURM, in its small way, is intended to help expand and formalize the methods by which new and conceptually variant models can be created, evaluated, and evolved. The ISL consists of an entirely new model of the liver. In and of itself, it may or may not prove useful in explanation or prediction of actual liver behavior. But when used and evolved side-by-side with data against which to validate and other trusted models with which to compare, it presents us with a vehicle for continual clarification of what may be going on within the liver without restricting our use of other models.

The ISL strives to replicate the experimental procedure that produced the experimental data sets. Obviously, there are overwhelming differences between any in silico model and an *in vitro* experimental biological system that no amount of discipline can reconcile. Examples of these in the ISL are the computational limitations that disallow (or dramatically increase the cost of) the simulation of moles of solute and a realistic population size of primary hepatic units and the cells that comprise them. However, because the ISL is validation-centric, what matters is whether or not the computation mimics the behavior of the real system to

the extent captured in the experimental data and within the other models being used. The ISL achieves this by focusing on the design, execution, and analysis of experiments. FURM is expanding modeling methods to cover and account for methods that have evolved and are used by nature to obtain the referent systems themselves. FURM, by positing methods for the creation, evolution, measurement, and selection of computational representations of functional units in biology, is doing just that.

ACKNOWLEDGMENTS

BIOGRAPHIES

Glen Ropella is a systems engineer with 13 years experience doing simulations, and for the past 7 years has worked in the domain of fine-grained Agent-Based Modeling. **C. Anthony Hunt** is a Professor of Biopharmaceutical Sciences and Bioengineering at the University of California, San Francisco; his research interests include Systems Biology, Computational Biology and the Pharmaceutical Sciences. Over the last ten years **Dev Nag** has conducted research in medical informatics, neuropsychology, and biological agent-based modeling.

REFERENCES

American Heritage Dictionary of the English Language: Fourth Edition. 2000. Houghton Mifflin Co., Boston, MA.

Beck, K. 2000. *Extreme Programming Explained.* Addison-Wesley Longman, Inc., Reading, MA.

Brenner, S. 2003. Keynote Address. *Biomedical Information Science and Technology Initiative (BISTI) 2003 Symposium* (Bethesda, MD, November 5-7).

Dewdney, A.K. 1989. "Computer Recreations: A Tinkertoy Computer that Plays Tic-Tac-Toe." *Scientific American* (October):120-123.

Di Paolo, E. A.; J. Noble; S. Bullock. 2000. "Simulation Models as Opaque Thought Experiments." *Proceedings of the Seventh International Conference on Artificial Life, Portland, OR.* MIT Press, Cambridge, MA, 497-506.

Ferber, D. 2004. "Synthetic Biology: Microbes Made to Order." *Science 303,* No. 5655:158-161.

Hung, D.Y.; P. Chang; M. Weiss; M.S. Roberts. 2001. "Structure-hepatic Disposition Relationships for Cationic Drugs in Isolated Perfused Rat Livers: Transmembrane Exchange and Cytoplasmic Binding Process." *J. Pharmacol. Exper. Therap 297*, No. 2:780-89.

Jacobson, I.; G. Booch; J. Rumbaugh. 1999. *The Unified Software Development Process.* Addison-Wesley Longman, Inc., Reading MA.

Kiczales, G.; J. Lamping; A. Mendhekar; C. Maeda; C.V. Lopes; J. Loingtier; J. Irwin. 1997. "Aspect-Oriented Programming." *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)* (Finland). *Springer-Verlag LNCS 1241.*

Krichmar, J.L. and G.M. Edelman. 2002. "Machine Psychology: Autonomous Behavior, Perceptual Categorization and Conditioning in a Brain-based Device." *Cerebral Cortex (August 12)*:818-830.

Lin, F.J. and Y. Lin. 1994. "A Building Block Approach to Detecting and Resolving Feature Interactions." In *Proceedings of the Second Int′l Workshop on Feature Interactions in Telecommunications Software Systems*, (Amsterdam, The Netherlands), 86-109.

McMullin, B. 2000. "John von Neumann and the Evolutionary Growth of Complexity: Looking Backward, Looking Forward." *Artificial Life 6*, No. 4: 347-61.

Metropolis, N.C. and S.M. Ulam. 1949. "The Monte-Carlo Method." *J. Amer. Stat. Assoc. 44*:335-341.

Roberts, M.S. and Y.G. Anissimov. 1999. "Modeling of Hepatic Elimination and Organ Distribution Kinetics with the Extended Convection-dispersion Model." *J. Pharmacokin. Biopharm. 27*, No. 4:343-382.

Ropella, G.; D. Nag; C.A. Hunt. 2003. "Similarity Measures for Automated Comparison of In Silico and In Vitro Experimental Results." *Proceedings of the 25th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC03)*, (Cancun, Mexico, September 17-21). IEEE, Piscataway, NJ, 2933-2936.

Ropella, G.E.P. and C.A. Hunt. 2003. "Prerequisites for Effective Experimentation in Computational Biology." *25th Annual Conference of the IEEE Engineering in Medicine and Biology Society* (Cancun, September 17-21).

Smith, B.C. 1996. *On the Origin of Objects.* MIT Press, Cambridge, MA.

Steels, L. and R. Brooks, eds. 1995. *The Artificial Life Route to Artificial Intelligence.* Lawrence Earlbaum Associates, Inc., Hillsdale, NJ.

Tarski, A. 1956. "The Concept of Truth in Formalized Languages." In *Logic, Semantics, Metamathematics.* Clarendon Press, Oxford.

Ulam, S.M. 1960. "Computing Machines as a Heuristic Aid." In *A Collection of Mathematical Problems.* Interscience Publishers, New York.

von Neumann, J. 1966. *Theory of Self-Reproducing Automata.* Edited and completed by Arthur W. Burks. Univ. of IL, Urbana, IL.

Walker, R.J. 2003. "Supporting Inconsistent World Views." *AOSD 2003 Workshop on Software Engineering Properties of Languages for Aspect Technologies (SPLAT 03)* (Boston, MA).

Wiley, B. 2000. *Essential System Requirements: A Practical Guide to Event-Driven Methods.* Addison-Wesley Longman, Inc., Reading MA.

Zeigler, B.P. 1984. *Multifacetted Modelling and Discrete Event Simulation.* Academic Press, New York.