

An Interactive Environment for Explanatory Biological Modeling

Pat Langley (langley@asu.edu)
Computing Science and Engineering
Arizona State University, Tempe, AZ 85287 USA

Abstract

In this paper, we describe an interactive environment for the representation, interpretation, and revision of explanatory biological models. We illustrate our approach on the systems biology of aging, a complex topic that involves many interacting components. We also report initial experiences with using this environment to codify an informal model of aging. We close by discussing related efforts and directions for future research.

Keywords: scientific models, qualitative reasoning, applied cognitive science

Introduction and Overview

There is general agreement that the explosive growth in biological data offers great opportunities but also poses major challenges. Although less widely recognized, the growing complexity of biological models that aim to account for these observations raises a host of other issues. Computational techniques hold promise for mitigating this complexity, but most responses have been driven by algorithmic concerns rather than the cognitive needs of scientists who must develop, interpret, and understand complex models. Biologists would benefit from new computational tools designed with scientific users in mind.

Many efforts in modern science aim to understand complex phenomena from a systems perspective. One important example comes from research on aging, with recent studies suggesting that senescence results from the interaction of many distinct but interconnected processes (Vijg & Campisi, 2008). Individual laboratories report experiments and propose hypotheses to explain them, but there has been little work on how they fit together. The systems biology movement has championed integrative science, but it has emphasized topics like gene regulation and left phenomena like aging understudied.

In this paper, we report an interactive computational framework designed to support modeling of this variety. Our approach relies on three distinct but mutually supportive ideas:

- formal representations of scientific knowledge that make contact with specific fields' terms and concepts;
- methods for reasoning over models cast in these formalisms that provide the same flexibility and draw the same conclusions as scientists;
- techniques that let researchers analyze and update these models in an incremental, cumulative manner.

In the next section, we discuss three computational challenges that these capabilities raise, after which we describe an interactive software environment that embod-

ies our responses. We illustrate the system's abilities with examples from the domain of aging, then report initial experiences with the environment. We conclude with a discussion of related work on scientific modeling, along with directions for additional research.

Some readers may question the relevance of our work to cognitive science. Of course, scientific reasoning has long been a topic of study within this community, but we will not claim our system reasons in precisely the same way as biologists. However, our approach is informed by results from cognitive science that constrain it in important ways. In particular, it borrows from research on qualitative mental models, which has proposed representations and reasoning methods that are consistent with knowledge about human cognition. A good analog comes from work on intelligent tutoring systems (e.g., VanLehn, 2006), which does not model the details of human tutors but takes lessons from them. We view our work on computational aids for biological modeling as another important instance of applied cognitive science.

Challenges in Scientific Modeling

As we have noted, the construction of complex scientific models raises three separate but interrelated challenges. Here we expand upon each of them in turn, placing constraints on the form our responses should take in developing an environment for biological modeling.

The overall aim of science is to produce knowledge, but the social nature of science requires the use of *communicable* formalisms that researchers can exchange and understand (Džeroski, Langley, & Todorovski, 2007). Thus, our first computational challenge involves selecting a communicable formalism for biological models. Over the past decade, computational researchers have proposed many notations for such models, but most utilize notations borrowed from other fields that have questionable relevance to traditional biological thinking. Research in biology generally, and on aging in particular, imposes two constraints on modeling formalism. One is that most accounts of phenomena are qualitative, not because researchers prefer them intrinsically, but because they enable useful claims even when lacking more precise information. A second feature is that biologists often move beyond simple predictive models to posit causal hypotheses or processes that underlie known phenomena.

Science also differs from some areas of inquiry by its concern with observations. However, biologists typically

desire more from their models than simple predictions; they prefer *explanations* that account for observations in terms of concepts and mechanisms they find familiar and plausible. Such explanatory reasoning is common in biology (Darden, 2006), but the growing complexity of models suggests that, without assistance, researchers will otherwise overlook important implications. Thus, a second computational challenge involves supporting reasoning over the communicable scientific formalisms just described. Methods for calculating results from numeric equations are well established, but automated reasoning over the qualitative models that dominate biology requires a different approach. One complication that arises in qualitative models is that two or more causal pathways can predict different relationships between variables. Another is that it can be difficult to reason qualitatively about how a system changes over time.

A third important feature of science is its cumulative character. Historians often focus on conceptual breakthroughs by individuals like Darwin, Pasteur, and Morgan, but the great majority of research involves filling in technical details rather than changing paradigms. This is especially true for biology and medicine, in which scientists devote considerable effort to piecing together complicated models with many interacting parts. Thus, our final computational challenge involves supporting the cumulative improvement of system-level models by biological researchers. A common response is to develop curated knowledge bases (e.g., Karp et al., 2000; Vastrik et al., 2007) that rely on centralized control by a few experts, but the field has also explored community-based approaches. Both require ways to update models incrementally as new knowledge becomes available.

An Interactive Modeling Environment

We have incorporated our responses to the above issues into an interactive software environment for biological modeling. We have implemented the initial system in Lisp and we have used it to formalize four compartments of Furber's (2009) network diagram of aging, which depicts in a graphical but informal way some well-supported hypotheses and phenomena from biogerontology. In this section, we report the environment's response to each of the challenges just described, using examples from aging to clarify its operation.

Representing Biological Models

Recall that our first computational challenge involves encoding explanatory models and presenting them in ways that biologists will understand. Let us review some key features of aging that hold implications for modeling these phenomena:

- Different effects of aging and age-related disease are localized in different portions of body. For instance, some age-linked changes occur in specific parts of the cell, such as the lysosome or the mitochondria.

- Some hypotheses about aging involve transient substances, such as enzymes and reactive oxygen species (ROS), whereas others involve far more stable entities like lipofuscin and mitochondrial mutations that accumulate over time.
- Empirical results generally take the form of qualitative relations between continuous variables. For instance, one robust finding involves a negative influence of caloric intake on lifespan in model organisms.
- Aging takes place over time, but its effects are primarily monotonic in character, with the values of variables increasing or decreasing consistently. For example, lipofuscin in the lysosome is generally observed to increase with chronological age.
- Empirical findings about aging come in two distinct varieties: uncontrolled observations about changes over time and results of controlled experiments that measure the effect of one variable on another.

Taken together, these observations provide both constraints on our approach to modeling aging processes and avenues for making the task more tractable.

Table 1 presents our reformulation of the lysosome compartment of Furber's network diagram. The initial 12 statements in (a) and (b) reflect the first two points above. They declare specific locations – the lysosome, the cytoplasm, and the cell that contains them – along with quantities that are measurable (at least in principle) in those locations. Some quantities refer to stable substances, such as junk protein, oxidized protein, and lipofuscin, which accumulate over time unless actively broken down, whereas others denote transient substances, like Fe, ROS, and lytic enzyme, which are reactive enough to be very short lived.

The table also includes a set of hypotheses (c) about how these quantities influence each other. One claim is that transient ROS increases with transient Fe within the lysosome, whereas another is that stable oxidized protein increases with transient ROS in the same location. Hypotheses may also relate quantities in distinct locations (e.g., that lipofuscin in the cytoplasm increases with damaged membrane in the lysosome). These hypotheses have a clear causal interpretation, in that they state how one variable will change when one alters another. However, although they link continuous quantities, the relations themselves are qualitative in character.

Of course, we should remember the purpose of hypotheses like those in Table 1 (c), which is to explain known empirical results and predict new ones. This in turn requires not only that we represent these empirical findings formally, but also that we distinguish them clearly from the hypotheses themselves. Table 1 (d) shows four facts about aging in the lysosome that illustrate our earlier point about two forms of empirical findings. The first two items clarify both the observational, nonexperimental character of many facts about

Table 1: Formalization of Furber’s (2009) lysosome model, including (a) locations, (b) stable and transient quantities in these locations, (c) hypothetical claims about causal influences between these quantities, and (d) empirical facts about relations between quantities.

(a)	location cell location lysosome in the cell location cytoplasm in the cell
(b)	stable junk protein in the lysosome and cytoplasm transient degradation rate in the lysosome transient Fe in the lysosome transient ROS in the lysosome stable oxidized protein in the lysosome stable lipofuscin in the lysosome and cytoplasm transient lytic enzyme in the lysosome stable damaged membrane in the lysosome transient H2O2 in the lysosome and cytoplasm
(c)	hypothesis junk protein decreases with degradation rate in the lysosome hypothesis junk protein in the lysosome increases with junk protein in the cytoplasm hypothesis Fe increases with junk protein in the lysosome hypothesis ROS increases with Fe in the lysosome hypothesis oxidized protein increases with ROS in the lysosome hypothesis lipofuscin increases with oxidized protein in the lysosome hypothesis degradation rate decreases with lipofuscin in the lysosome hypothesis lytic enzyme decreases with lipofuscin in the lysosome hypothesis ROS increases with lipofuscin in the lysosome hypothesis damaged membrane increases with ROS in the lysosome hypothesis lipofuscin in the cytoplasm increases with damaged membrane in the lysosome hypothesis H2O2 in the lysosome increases with H2O2 in the cytoplasm
(d)	fact lipofuscin in the lysosome increases with time fact membrane damage in the lysosome increases with time fact lytic enzyme decreases with ROS in the lysosome fact H2O2 does not change with ROS in the lysosome

aging and also their monotonic nature. These explicitly mention time as a variable, which the model hypotheses do not. The other two facts reflect (plausible) results of experimental studies that measure the effect of one quantity’s variation on another. The first states that lytic enzyme decreases with ROS in the lysosome. The second states that H2O2 does not vary with of ROS. Such negative results place constraints on models, although hypotheses may contain only positive causal relations.

This notation meets two of the criteria given earlier. It supports qualitative models that nevertheless relate quantitative variables of the type that biologists typically measure, and the hypotheses that make up models




Figure 1: A graphical visualization of the qualitative lysosome model from Table 1, with plus (+) on an arrow denoting that one quantity increases with another and with minus (−) denoting a decreasing relationship.

have a clear causal interpretation. The formalism also lends itself to graphical display, with quantities shown in locations where they occur and with arrows depicting direct causal influences between these variables. Figure 1 shows a graphical version of the lysosome model from Table 1, with the empirical facts omitted. Our implemented system does not yet generate such graphs automatically, but adding this ability should not be difficult.

In addition, our notation lets users specify places, quantities, hypotheses, and empirical facts in constrained English, which we believe will make it more accessible to biologists who are uncomfortable with traditional computer languages. Yet models stated in this notation are well defined and unambiguous about their claims, making them just as formal as ones stated in the more arcane languages typically proposed in computational biology. This also distinguishes our approach from work on qualitative reasoning in cognitive science and AI (e.g., Bredeweg et al., 2007; Forbus, 1984), which has influenced our approach to biological modeling.

Reasoning over Biological Models

Our second computational challenge involves interpreting a given model to account for known phenomena. Scientists regularly engage in such reasoning, but with complex models they can easily overlook some conclusions and incorrectly infer others (e.g., Feldman et al., 1989). Thus, automatically determining a model’s implications should be a key part of our scientific modeling environment. Good models should explain known phenomena and predict new ones accurately, while phenomena place constraints on model content. Our framework’s formal statement of hypotheses and empirical results has another advantage: it lets one answer questions about how one quantity should affect another and predict the outcomes of thought experiments.

We can clarify this ability by introducing the notion of a *query* about how two quantities are related. This takes the same form as an empirical finding except that it does not state the direction in which one variable influences another or indeed whether an influence occurs at all. Thus, given the hypotheses in Table 1, we might ask “Does lipofuscin in the cytoplasm vary with Fe in the lysosome?” or “Does ROS in the lysosome vary with time?” The first asks a question about how changes to one quantity in a controlled experiment affect another; the second asks how a given quantity changes over time. The reasoning task is simplified by our assumption that effects are monotonic in character, giving behavior that one can describe in terms of a single qualitative state. This differs from much work on qualitative reasoning, which deals with trajectories of such states over time (e.g., Bredeweg et al., 2007; Forbus, 1984).

Because hypotheses take a form similar to facts, we can utilize a relatively straightforward chaining procedure to answer queries. To handle a question about how dependent variable Y varies with independent variable X, other things being equal, one simply finds a causal pathway, typically through other quantities, that starts with Y and ends with X. If no such path exists, then one can conclude that changes to X do not produce changes in Y. If there is such path, then one must still predict the direction of the effect. Briefly, if the path contains an even number of ‘decreases’ links, then one predicts that Y increases with X; otherwise one predicts that it decreases. For example, the model in Figure 1 lets one conclude that lytic enzyme will decrease with ROS. The justification for this strategy is simple: each ‘decreases’ link reverses the direction or sign of the path’s overall influence, so that an even number of them cancel out.

One complication arises when multiple paths from Y to X make different predictions. Without knowing the functional forms and parameters that produce each causal link, one cannot determine the exact effects of alternative pathways. Given the modeling framework as we have described it, in such cases one can only state that the hypotheses make contradictory predictions. However, we can extend the formalism in a simple way that lets it express another type of hypothesis that biologists regularly make: that the effect of one causal pathway dominates that of another. This requires a way to specify paths between two quantities and note which has the greater or dominating effect. Once included, such dominance relations let a qualitative causal model make unambiguous predictions about how one quantity varies with another, despite its abstract character.

Reasoning about how quantities change over time requires a slightly different approach. We assume that any exogenous variables not influenced by other quantities take on constant positive values. One can then infer the effect of such an exogenous quantity on another variable

downstream by finding pathways that connect them and combining the influences on their causal links. One can conclude that ‘stable’ quantities occurring downstream will increase or decrease over time, depending on their relation to the exogenous term. We can treat causal loops between two variables as special cases of conflicting paths in which a variable influences itself, again provided we specify which path is dominant.

Taken together, these computational mechanisms respond to a number of the issues raised above. They let our biological models move beyond inert structures to become interpretable ‘programs’ one can use to answer directed queries and make predictions about empirical relations. They also support reasoning about the effects of both controlled manipulation and the passage of time. As we will see shortly, the system can also explain the reasons for its conclusions. Computational aids of this sort should let biologists derive the implications of system-level models of aging that are more complex than ones they can handle without assistance.

Interactive Aids for Model Improvement

Our third computational challenge involves the incremental revision of models to bring them into closer alignment with known phenomena. This depends on the ability to represent such models and reason over them, but it must go beyond to identify portions of models that are problematic and modify them in response. Although there has been some research on automated model revision (e.g., Mahidadia & Compton, 2001), we have chosen to rely on interactive revision under user control. To this end, the system includes a number of commands through which users can update the knowledge base. These are currently available only through a textual interface, but we also plan to embed them in a graphical environment.

Naturally, the most basic commands includes ones for adding new model elements. The user can introduce new locations, quantities, hypotheses, and empirical facts by entering this content in the same format as shown in Table 1. The modularity of the modeling formalism, and its constrained English syntax, make these steps simple to carry out. The environment also includes a display command that presents the user with all elements in the current model or only those of a specified type. These commands provide the basic functionality needed for the cumulative improvement of causal biological models.

However, the system also provides users with additional details about the model’s behavior that can inform their revisions. In addition to answering specific queries like “Does ROS in the lysosome vary with time?”, users can also ask the environment to compare the current model’s predictions to known phenomena. When these predictions disagree with the empirical facts, the user can also ask the system to explain its reasoning. For each explanation, it presents the causal chain between two quantities that, taken together, predicted a partic-

ular outcome. Exceptions occur when the model incorrectly predicts no effect because no causal chain exists or makes an ambiguous prediction when two paths conflict and the user has not specified one as dominant.

The ability to inspect not only predictions but the reasoning behind them provides important insights about a model's strengths and weaknesses. If the model fails to match one or more empirical facts, explanations may reveal the source of the problem and ways to fix it. The user can remedy such situations in two basic ways – by adding new hypotheses, as described above, and by removing existing hypotheses. However, because the impact of deleting an element may be unclear in advance, the environment also lets users disable a model element without removing it entirely, as well as enable it later if that seems desirable. Taken together, these commands provide basic support for the incremental improvement of models, which will continue to be needed as new phenomena become available and demand explanation.

Initial Experiences with the Environment

We selected the systems biology of aging as our initial application domain because it was gaining increased attention within biology and because John Furber (2009) had already developed a network diagram that summarized many hypotheses and phenomena in this complex field. Repeated discussions with Furber let us convert his informal statements into our modeling notation.

We have focused our efforts on four compartments of Furber's diagram. These involve the dysfunction of lysosomes due to the accumulation of indigestible aggregates known as lipofuscin, the degeneration of mitochondrial energy production in the cell as the result of mutations, the shortening of telomeres and decline in Lon protease mRNA over time in the cell nucleus, and the crosslinking of proteins in the extracellular matrix. The lysosomal model, already seen in Table 1 and Figure 1, incorporated three places, nine quantities, and 12 hypotheses. The mitochondrial model included three places, nine quantities, and ten hypotheses, while the nuclear and extracellular models have similar complexities.

Naturally, translation of content from the informal diagram into our logical notation required some care and effort, with certain representational issues becoming apparent only along the way. Interactions with Furber clarified his intentions and usually determined how to proceed. Once we had the initial translation complete, we used the environment to detect and correct problems with these models, much as we intend its use by scientists. Running the reasoning mechanism over these models revealed a number of errors, some in our encoding of Furber's chart but also a few ambiguities in the original aging diagram itself. Formalization of the aging model, combined with the environment's reasoning methods, led to repair of these problems.

Related Work on Scientific Modeling

Our approach to interactive biological modeling borrows ideas from three distinct traditions, but combines them in new ways to produce novel capabilities. The computational biology community has pursued a number of projects that support Web-based access to biological knowledge. For instance, KEGG (Kanehisa, 1997), Reactome (Vastrik et al., 2007), and Metacyc (Karp et al., 2000) let their users explore biological content that curators have extracted from the literature, but they have only limited abilities to reason over their knowledge.

Some other biological modeling efforts come closer to our framework. For example, Genepath (Zupan et al., 2003) offers a Web-based environment that lets users enter qualitative results from genetics experiments and knowledge about gene regulation, but the model construction process is entirely automated. JustAid (Mahidadia & Compton, 2001) supports iterative revision of qualitative causal models, with the system proposing changes but the user selecting which to implement. Racunas et al.'s (2004) HyBrow supports interactive creation of qualitative models and checks their consistency with logical reasoning, but our system provides a more general treatment of explanatory biological models.

Of course, we have also been strongly influenced by research on mental models in cognitive science, especially work on qualitative reasoning and simulation (e.g., Forbus, 1984). Our approach shares some key ideas, especially that models involve qualitative causal relations among continuous variables. One difference is our assumption that behavior is monotonic over time, which simplifies reasoning considerably. Another distinction is our willingness to resolve ambiguity by specifying that one path dominates another. A third lies in our emphasis on predicting relations between pairs of quantities, rather than on model simulation. Our incorporation of qualitative models into an interactive modeling environment is not new. Bredeweg et al.'s (2007) GARP lets users construct qualitative models manually and simulate their behavior, although it focuses on ecology rather than biology, it uses a more complex process ontology, and it does not emphasize incremental revision.

Directions for Future Research

Although our modeling environment shows considerable promise, we need to extend the framework along a number of fronts. Clearly, our first step should be to embed the existing abilities in a graphical interface. This would let users visualize models in a manner similar to Figure 1, but it would also use this display to support query answering, prediction, and explanation, each of which have natural visual analogs. The environment would include templates for creating new locations, quantities, hypotheses, and empirical facts, for disabling and enabling model elements, and for copying and editing entire

models. These features would not change the environment's basic functionality, but they would make it more accessible to many biologists.

We should also expand the representational abilities of the modeling framework. One extension would enable grouping a set of causal links into a process, much as in Forbus' (1984) qualitative process theory. This would let a graphical interface hide model details until a user asks to see individual connections. Another augmentation would allow contextual conditions on causal links that specify the tissues and organisms in which they occur. If queries included similar conditions, then the reasoning system would collect relevant connections to create query-specific models for use in drawing conclusions. Finally, we should explore ways to move beyond the framework's strict assumption of monotonic behavior. One response would involve adding quantitative conditions to causal links and dominance relations that specify when they hold, with the reasoner collecting relevant model elements to make predictions for a specific situation.

Concluding Remarks

In this paper, we reported an interactive approach to the representation, interpretation, and revision of scientific models. Our environment encodes models as sets of qualitative causal influences that relates quantities in particular location, and its reasoning methods answer queries, make predictions, and explain its conclusions. Users can interactively invoke these abilities, which should help them understand a model's behavior and improve it over time. We have carried out initial tests on cellular models of aging, using the environment's interactive character to identify problems in these models and repair them.

Although our approach draws on ideas developed in earlier work, it combines them in novel ways to support three key facets of the scientific enterprise: the formal representation of knowledge and hypotheses, relating that knowledge to observations through explicit reasoning, and the incremental development of knowledge over time. Many projects that formalize biological knowledge have focused on inert structures, rather than offering aids for reasoning over complex models, and most techniques for codifying knowledge rely on curators, rather than giving scientists tools to make their own changes. We believe our interactive environment offers a promising approach that addresses these issues in ways that biologists will find accessible and useful.

Acknowledgements

This research was supported in part by Grant CAA 0113-07 from Science Foundation Arizona and in part by Arizona State University. We thank John Furber for providing feedback about his aging network diagram, along with Durga Bidaye, Rick Chimera, Juraž Dzifcak, Seungchan Kim, Stephen Racunas, David Stracuzzi, and Michael Verdicchio for early contributions to the project.

References

- Bredeweg, B., Bouwer, A., Jellema, J., Bertels, D., Linnebank, F., & Liem, J. (2007). Garp3 - A new workbench for qualitative reasoning and modelling. *Proceedings of the Fourth International Conference on Knowledge Capture* (pp. 183–184), Whistler, BC.
- Bridewell, W. & Langley, P. (2010). Two kinds of knowledge in scientific discovery. *Topics in Cognitive Science*, 2, 36–52.
- Darden, L. (2006). *Reasoning in biological discoveries*. New York: Cambridge University Press.
- Džeroski, S., Langley, P., & Todorovski, L. (2007). Computational discovery of scientific knowledge. In S. Džeroski & L. Todorovski (Eds.), *Computational discovery of communicable scientific knowledge*. Berlin.
- Feldman, B. Z., Compton, P. J., & Smythe, G. A. (1989). Hypothesis testing: An appropriate task for knowledge-based systems. *Proceedings of the Fourth Knowledge Acquisition for Knowledge-based Systems Workshop*. Banff, Canada, October 1989.
- Forbus, K. D. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85–168.
- Furber, J. (2009). Systems biology of human aging: Network model of biochemical and physiological interactions in human senescence. <http://www.legendarypharma.com/chartbg.html>.
- Kanehisa, M. (1997). A database for post-genome analysis. *Trends in Genetics*, 13, 375–376.
- Karp, P., Riley, M., Saier, M., Paulsen, I., Paley, S., & Pellegrini-Toole, A. (2000). The EcoCyc and MetaCyc databases. *Nucleic Acids Research*, 28, 56–59.
- Mahidadia, A., & Compton, P. (2001). Assisting model-discovery in neuroendocrinology. *Proceedings of the Fourth International Conference on Discovery Science* (pp. 214–227). Washington, DC: Springer.
- Racunas, S., Shah, N., Albert, I., & Fedoroff, N. (2004). HyBrow: A prototype system for computer-aided hypothesis evaluation. *Bioinformatics*, 20, i257–264.
- VanLehn, K. (2006). The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, 16, 227–265.
- Vastrik, I., D'Eustachio, P., Schmidt, E., Joshi-Tope, G., Gopinath, G., Croft, D., de Bono, B., Gillespie, M., Jassal, B., Lewis, S., Matthews, L., Wu, G., Birney, E., & Stein, L. (2007). Reactome: A knowledge base of biologic pathways and processes. *Genome Biology*, 8, R39.
- Vijg, J., & Campisi, J. (2008). Puzzles, promises and a cure for ageing. *Nature*, 454, 1065–1071.
- Zupan, B., Bratko, I., Demsar, J., Juvan, P., Halter, J. A., Kuspa, A., & Shaulsky, G. (2003). GenePath: A system for automated construction of genetic networks from mutant data. *Bioinformatics*, 19, 383–389.