AAAI Technical Report FS-12-03
Discovery Informatics: The Role of AI Research in Innovating Scientific Processes

A Web-Based Environment for Explanatory Biological Modeling

Pat Langley and Glen Hunt
(langley@asu.edu) (glen.hunt@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, and discuss our experiences using this environment to codify an informal model of aging. We close by discussing related efforts and directions for future research.

Introduction and Overview

There is general agreement that the explosive growth in biological data offers great opportunities but also poses major challenges for many fields. 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 efforts in this arena have been driven by algorithmic concerns rather than the cognitive needs of scientists who must develop complex models, interpret them, and understand their behavior. Researchers would benefit from computational tools designed with their needs in mind.

Many efforts in modern science aim at understanding complex phenomena that require a systems perspective. One important example comes from research on aging in humans and other organisms, 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, Web-based computational framework designed to support modeling of this variety. We illustrate the system's abilities with examples from the aging domain, then report initial experiences with the environment. We conclude with discussions of related work on scientific modeling and directions for additional research.

Challenges in Scientific Modeling

The construction of complex scientific models raises three distinct 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.

Communicable Scientific Formalisms

The overall aim of science is to produce knowledge. Other areas of human endeavor share this goal, but one way that science differs is its emphasis on formalized statements of this content. However, the social nature of science imposes an additional constraint: it must utilize *communicable* formalisms that researchers can exchange and understand, even over great distances (Džeroski, Langley, & Todorovski, 2007). Different fields have developed distinct notations to convey their knowledge, each well suited to its community's needs.

Thus, the first computational challenge we must address is to select a communicable formalism for biological models. Over the past decade, many notations for formal biological modeling have been proposed, but most involve notations borrowed from other fields. These have included differential equations, Bayesian networks, and Boolean networks, all of which have questionable relevance to traditional biological thinking. We hold that research in biology generally, and on aging in particular, imposes three constraints on modeling formalism. One is that most accounts of phenomena are qualitative in character, not because researchers prefer them intrinsically, but because they enable useful claims even when lacking more precise information. Another factor is that most biologists prefer graphical or diagrammatic notations of model content. A third feature is that biologists often attempt to move beyond simple predictive models to posit causal hypotheses or processes that underlie known phenomena. A successful formalism should respond to all of these factors.

Explainable Scientific Reasoning

Science also differs from some areas of inquiry by its concern with observations. For a scientific model to be useful, it must make some connection to data or phenomena. This in turn requires some form of reasoning that leads from the

model to predictions that, ideally, are consistent with observations. However, biologists typically desire more 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. Different types of formalism and data depend on different kinds of reasoning. Some of these, like methods for calculating results from numeric equations, are well established. Automated reasoning over the qualitative models that dominate biology requires a different approach, with techniques from logic being natural candidates. One complication that arises in qualitative models is that two or more causal pathways can predict different relationships between variables. Another is that it is difficult to reason qualitatively about how a system changes over time. Successful computational aids must address these issues.

Cumulative Improvement of Models

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. The fact that this cumulative work happens in a distributed community introduces additional complexity to the enterprise. Some fields have managed reasonably well with traditional tools of exchange like refereed journals, but the systems biology of aging would benefit from more advanced technologies.

Thus, our final computational challenge involves supporting 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. For this reason, curated methods do not scale well and can take considerable time to incorporate recent findings. We favor a more open approach in which many different researchers extend and revise a widely known model and, if they decide the results are worth sharing, make them available to others.

An Interactive Modeling Environment

We have incorporated our responses to the above issues into a Web-based software environment for biological modeling. We have used it formalize two 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 and Visualizing 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 reactive oxygen species (ROS), whereas others involve 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.
- 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 place strong constraints on our approach to modeling aging processes.

Figure 1 presents our reformulation of the lysomone compartment of Furber's network diagram. The initial 15 statements in the textual display on the left 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, like junk protein, oxidized protein, and lipofuscin, whereas others denote transient substances, like Fe, ROS, and lytic enzyme. Locations and quantities are stored internally as simple frames, with the values for some slots (e.g., location of a H2O2 instance) pointing to others (e.g., the lysosome).

The textual display in the figure also includes a set of hypotheses 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 changes when another is altered. However, although they link continuous quantities, the relations themselves are qualitative in character. Hypotheses are also encoded as frames, with some slots (e.g., cause and effect) referring to instances of defined quantities.

Remember that the purpose of hypotheses 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

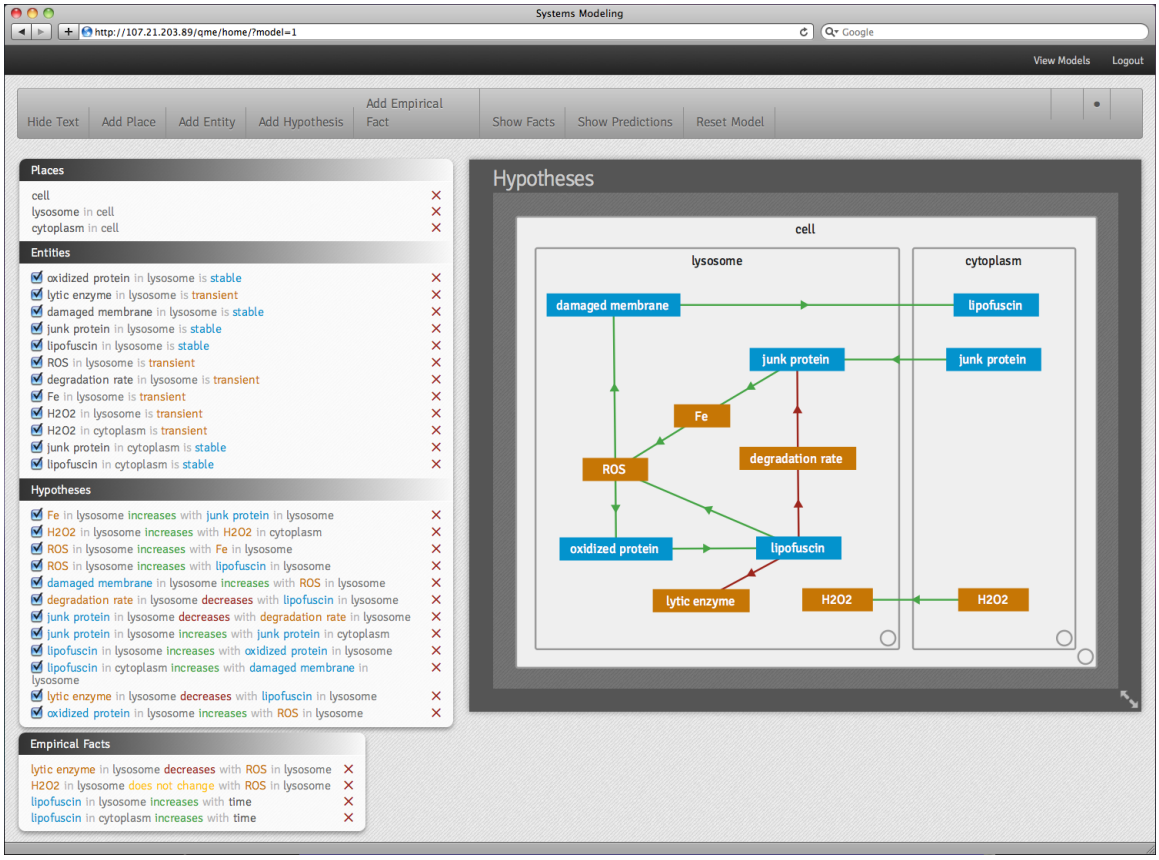

Figure 1: A screen shot of the Web-based modeling environment. The textual display (left) items into places, stable and transient quantities that occur in these locations, hypothesized causal influences that relate quantities, and empirical facts. The graphical display (right) depicts places as boxes, quantities as nodes within those boxes, and causal hypotheses as links connecting them.

hypotheses themselves. The final four items on the left of the figure illustrate our earlier point about forms of empirical findings. The first two items both reflect the observational, nonexperimental character of many facts about aging. These explicitly mention time as a variable, which the model hypotheses do not. The other two facts reflect the results of experimental studies that measure the effect on one quantity when another is varied. 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 biologists typically measure, and the hypotheses that make up models have a clear causal interpretation. The formalism also lends itself to graphical display, as shown in the right side of Figure 1. Here places are depicted as boxes, quantities appear in their locations, and arrows denote direct causal influences between these variables, with green for positive influences and red for negative ones.

Users can also display empirical facts in a graphical format. This material takes the place of the textual display, showing the observed relations in the same layout as as hypotheses that make up the model. The difference is that the latter presents direct connections between quantities intended to explain the indirect connections observed empirically.

Reasoning over Biological Models

The second computational challenge involves interpreting a given model to explain known phenomena. Scientists regularly engage in such reasoning, but with complex models they can easily overlook some conclusions and incorrectly infer others. 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, while phenomena place constraints on model content.

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 of influence or indeed whether an influence occurs at all. Thus, given the hypotheses in Figure 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 quantity changes over time.

Because hypotheses take a form similar to facts, we can utilize a relatively straightforward chaining procedure to answer such 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 if the path contains an even number of 'decreases' links, one predicts that Y increases with X; otherwise one predicts that it decreases. For example, the model in Figure 1 implies that lytic enzyme will decrease with ROS.

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. The current modeling 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 make all the time: that the effect of one causal path dominates that of another.[1] This requires a way to specify which path between two quantities has the greater effect. Such dominance relations let an abstract, qualitative causal model make unambiguous predictions about how one quantity varies with another.

Reasoning about how quantities change over time, rather than as the result of experimental control, requires a slightly different approach. We assume that any exogenous variables not influenced by other quantities is constant. We can 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. We can infer that 'stable' quantities occurring downstream will increase or decrease over time, depending on their relation to the exogenous term. For instance, the model in Figure 1 implies lipofuscin will increase monotonically in this way.

Taken together, these computational mechanisms transform our biological models from inert structures into become interpretable 'programs' one can use to make predictions about empirical relations and to explain the reasons for these conclusions. They also support reasoning about both controlled experiments and observed effects over time. Computational aids of this sort should let biologists derive the implications of system-level models of aging 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 users must also 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 actions through which users can update the knowledge base. These are available only in the textual frame, but we also plan to support analogous graphical commands.

Naturally, the most basic commands includes ones for adding new model elements. Users can introduce new locations, quantities, causal hypotheses, and empirical facts by entering this content in the same format as shown in Figure 1. The modularity of the modeling formalism and the simple structure of each element make these steps easy to carry out.

Entry involves selecting a type of element to add and then completing a number of pull-down menus to describe the new entry. For example, entering a new hypothesis involves clicking the 'add hypothesis' button and then selecting the influenced quantity, the influencing quantity, and the direction (increases or decreases) from the menus. Users can only select quantities that were entered earlier, thus ensuring that hypotheses only relate known variables. Similar interactive commands enable the introduction of new places, quantities (which specify a name, place, and whether stable or transient), and empirical facts. Users can also remove items they feel are no longer necessary by clicking a 'delete' button to the right of the item in question. Removing quantities also leads to removal of all hypotheses and facts that refer to them, while removing places leads to deletion of all quantities that reside in those locations.

Together, these commands provide the basic functionality needed to construct causal biological models, but model development is an incremental process that aims to bring hypotheses into better alignment with empirical evidence. For this reason, the environment also lets users examine the current model's predictions and their relation to known phenomena. When users click on the 'show predictions' button, the system derives the predictions associated with each empirical fact and presents them in the textual or the graphical display, whichever is currently on the screen. The textual frame shows predictions directly below their corresponding facts, with a marker indicating whether each pair agrees or disagrees. The graphical frame displays predictions as thicker lines overlaid on the arrows or trapezoidal boxes that denote empirical facts. The same colors mean the prediction agrees with the observation, while different colors indicate a problem that needs attention.

Of course, before users attempt to revise their models, they should understand the reasons for faulty predictions. To this end, the environment lets them inspect the explanations associated with each one.[2] We have not yet implemented this ability for the textual frame, but, when users click on a prediction line in the graphical frame, explanations are shown in the hypothesis frame. If only one explanation exists for the prediction, then the system highlights the causal links involved and places the others in the background. If multiple explanations exist, then the environment shows a number for each alternative and, when the user clicks on a given number, highlights only the causal links for that ac-

[1] Another complication involves models in which causal loops occur between two quantities, which we can handle as special cases of models with multiple paths.

[2] Users can also examine explanations for correct predictions, but these are less useful in improving the model.

count. This mechanism also lets users understand cases in which the model makes ambiguous predictions.

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 facilities provide users with the information they need to incrementally refine their qualitative causal models to bring them into alignment with empirical findings.

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 focused our efforts on two compartments of Furber's diagram, one involving the dysfunction of lysosomes due to the accumulation of indigestible aggregates known as lipofuscin, and another on the degeneration of mitochondrial energy production in the cell as the result of mutations. The lysosomal submodel, already seen in Figure 1, incorporated three places, 12 quantities, and 12 hypotheses. The mitochondrial submodel included three places, nine quantities, and 11 hypotheses, with the two models having little overlap.

Naturally, translation of content from the informal diagram into our logical notation required some care and effort, with some 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 submodels, much as we intend its use by scientists. Running the reasoning mechanism over these submodels 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 users explore biological content that curators have extracted from the literature, but they do not reason over the knowledge or let users modify it.

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 model construction process is entirely automated. JustAid (Mahidadia & Compton, 2001) supports iterative revision of qualitative causal models, with the system proposing changes but users 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.

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 is our willingness to resolve ambiguity by specifying that one path dominates another. 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 and uses a more complex process ontology.

Directions for Future Research

Although our modeling environment shows considerable promise, we need to extend the framework along a number of fronts. Our first step should be to introduce graphical versions of all commands for adding and revising content, which should make the environment more accessible to many biologists. We should also support 'thought experiments' by asking the system to make predictions in the absence of empirical results. In addition, users would benefit from the ability to copy and edit entire models, as well as overlay alternative models to reveal their similarities and differences.

Expanding the representational abilities of the modeling framework should be another priority. One extension would enable grouping a set of causal links into a process that would let the interface hide model details until users ask to see individual connections. Another augmentation would allow contextual conditions on causal links that specify the tissues and organisms in which they occur. Given similar conditions on queries, the reasoner would collect relevant connections to create query-specific models for use in drawing conclusions.

We should also explore ways to move beyond two assumptions – monotonic behavior and pathway dominance – that we introduced for practical reasons. One response to both would involve adding quantitative conditions with arithmetic formulas to causal links and dominance relations that specify when they hold. The reasoner would then collect relevant model elements to make qualitative predictions

based on quantitative measurements. Also, because pathway dominance relations violate the compositionality of traditional qualitative models, we should develop means for decomposing them into fragments by removing shared links and identifying subpathways responsible for the dominance.

We are collaborating with experts in biogerontology who intend to populate the modeling environment with content from the literature about human senescence. The aim is to support thought experiments about interventions that could delay or reverse aspects of aging, which in turn could lead to empirical studies that test these predictions. We anticipate that this approach will suggest novel treatments that are nonobvious to biologists due to long causal chains, which our system can find without difficulty. We believe that the envronment will prove equally useful in modeling other areas of biology, such as cancer and Alzheimer's disease, that involve monotonic changes over time.

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 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 in the lysosome and mitochondria, using the interactive character of the environment 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 DaimlerBenz Research and Technology Center. We thank John Furber for providing feedback about his aging network diagram and Stephen Racunas for early contributions to the project. Portions of this paper have appeared in *Proceedings of the Thirty-Second Annual Meeting of the Cognitive Science Society*.

References

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.

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.

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., Kuspa, A., & Shaulsky, G. (2003). GenePath: A system for automated construction of genetic networks from mutant data. *Bioinformatics*, *19*, 383–389.