5

On building reliable pictures with unreliable data:[1] An evolutionary and developmental coda for the new systems biology

William C. Wimsatt

SUMMARY

The new systems biology (NSB) is a cluster of methodological approaches for the analysis of dynamical behavior in networks that sits at the confluence of a number of disciplines. They are all now experiencing massive increases in qualitatively new actual and potential interactions driven by the data explosions in genomics and proteomics. I urge that developmental and evolutionary perspectives provide particularly useful tools for the analysis of life systems in the NSB. Their character as systems that must develop and evolve means that they possess certain properties, in particular evolvability, genetic and environmental robustness, and differential generative entrenchment for their parts. These properties are themselves very general and robust. They also possess virtues that help to ameliorate a problem of rapidly growing magnitude in the analysis of complex living systems: that the data produced by high-throughput methods ('gene chips') have very high error rates. Some of these errors are undoubtedly products of a technology that is new, noisy, and needs further tuning. Some are almost certainly systematic and, by recognizing this, remediable. Knowledge of the general kinds of systems we are dealing with can help with both.

[1] This title is an homage to the famous paper of a similar name by John von Neumann, published posthumously 50 years ago, in 1956. Von Neumann's (1956) paper pioneered the consideration of reliability, in both organic and computer design.

Systems Biology
F.C. Boogerd, F.J. Bruggeman, J.-H.S. Hofmeyr and H.V. Westerhoff (Editors)
Isbn: 978-0-444-52085-2

1. INTRODUCTION

These are exciting times for biology, on multiple fronts. A number of disciplines intersect and are enlightened by the explosive growth of new knowledge in detail and at the system level about genetic and biochemical interactions. New perspectives on the evolution, phylogeny, development, and organization of complex adaptive systems emerge as we learn more about these interacting systems in development at the biochemical, cellular, and multicellular levels affecting differentiation and compartmentalization. The new systems biology (NSB) is riding an expanding wave as we found and rename departments in its name.

We also seem to be reverting spontaneously to talk that was more common in the heyday of 'systems theory' and cybernetics in biology, from the late 1950s to the early 1970s. This reversion is a product of the kinds of knowledge we are gaining. It was more schematic and promissory then; now, while still schematic in many places, it is increasingly richly empirically based and detailed. The timeliness of much of this history is explored elsewhere in this volume by Evelyn Fox Keller.[2] And increasing use of the cybernetic vocabulary comes from both macro- and microdirections: Thus Wallace Arthur (1997) and Eric Davidson (2001) both make rich use of such language and 'wiring diagrams' of interactions between and among genes and their products. Arthur, a self-retooled population biologist, became intrigued by the rich complexities of development. His interest is morphological but reaches down to detailed gene-control interactions relevant to morphological expression.[3] Davidson, a pioneer in the study of gene control from the late 1960s and fairly speaking, a new systems biologist before there was a NSB, has a focus that is more 'bottom-up', analyzing, and articulating gene-control networks and cascades to extrapolate to an overall, developmental architecture (Davidson & Erwin, 2006). The return to cybernetic language is not surprising. In the last decade, we have analyzed 'genetic wiring diagrams' of increasing complexity and scope (Davidson, 2006).

---

[2] Evelyn and I are both historically well placed to remember it though as a biophysicist working on development she was a participant, while I looked on enviously, by then as a philosopher. My connections came through Frank Rosenblatt's broad ranging course in the Fall of 1964 (titled 'Brain Models and mechanisms', but it was really on adaptive systems more generally and many of the readings showed the influence of cybernetics and systems theory). Of all that I read, probably Kacser (1957) came closer to representing the spirit of the NSB. It was not well known, but in many ways the spirit of the NSB was paradigmatically anticipated.

[3] Another example with another approach is provided by Stuart Newman, a theoretical chemist by training, and an evolutionary morphologist Gerd Müller. They undertake systematic exploration of morphological possibilities for cellular constructions and their connections with the underlying chemistry (Newman and Müller 2000; Müller and Newman 2003). Their approach seeks generic constraints on possible modes for assembly of cells into larger morphological structures. In the generality it seeks, it is a methodology more reminiscent of bottom-up approached from physics, but practiced on top-down objects and phenomena.

2. THE NEW SYSTEMS BIOLOGY AND EVO-DEVO

Is evo-devo irrelevant? To some in the NSB, it simply indicates another domain that NSB can do without. I disagree. So would Eric Davidson and Evelyn Fox Keller. So, I expect would Carl Woese, a pioneer in the elucidation of the genetic code and in the phylogeny of very early life. Those who believe they have no need of evolutionary and developmental perspectives have perhaps been misled by a common kind of stereotyping of these disciplines.

Many practitioners of evolutionary developmental biology themselves feel that 'evo-devo' lumps diverse practices and perspectives, making them appear more monolithic than they feel. To some, it seems too skewed towards developmental genetics as opposed to higher levels of organization such as either morphology (Love, 2003) or ecology (Gilbert, 2001). Others would worry about the emphasis on long time scales and correlative emphasis on typological conceptions of species away from populational variability (Raff, 1996, p. 21). In this stereotypic image, population genetics, while dynamic, is limited to models of genetic change in terms of selection coefficients and gene frequencies that abstract away from physiology and phenotypic organization completely, and the evolutionary biology of macroevolution is just descriptive and not predictive. No wonder it seems irrelevant. This stereotyping is inevitable for a new discipline that articulates so many prior separate areas, and in which most of the practitioners of one subarea are amateur consumers of most of the others. But unfortunately it masks a great deal of relevant work.

In fact the confluence of developmental genetics with systems approaches and a more macroscopic developmental biology, systematics, and comparative studies has created a new hybrid discipline in which population genetics can move towards more detailed dynamical models of the phenotype. It is doing so both synchronically and through its developmental history, and in which systematics and phylogeny are again used to provide important clues to the organization of development and the course of evolution, predictively as well as descriptively, as they did in the nineteenth century. At the hands of researchers like systematist–geneticist–systems biologist Carl Woese, they are redoing the history of the early origins of life and uncovering surprising things about its nature (Woese, 2004). These include the initially shocking claims of the endosymbiotic origins of eucaryotes (Margulis, 1971) and its subsequent elaboration to discover widespread, now entrenched, symbioses. Morowitz's (1992) claim that large chunks of metabolism represent preserved chunks of earlier biotic environments does have things to say about the origins and nature of life, as well as about its evolution. More recently, Woese's (1998) hypotheses that there was an early stage in which the ancestors of already distinct lineages interchanged genes far more readily than later after the development of mitosis solved a remaining puzzle for systematists with a set of physiological and evolutionary proposals.

These suggest investigations that could be done by systems biologists concerning the origins and plausible evolutionary order of the different complex elements of mitosis. Hypotheses about the origins of life have always had to pass muster on biochemical grounds and plausibility. All these have implications for the broader architecture of living systems within which NSB works. Not everything in either evolutionary biology or developmental biology is relevant to NSB, but some of it is already part of NSB, and other parts will become increasingly relevant over time as we learn to better relate processes, acting on different time and size scales.

We should not fear that evolutionary and developmental biology will simply swallow systems biology, because systems biology is characterized by its approach as much as by its subject matter. Nor must the aim of the NSB be to serve developmental and evolutionary biology, any more than it might be to serve, e.g., oncology or epidemiology. Evolutionary biology, developmental biology, genetics, cell physiology, and biochemistry are using converging methodologies on the common stage of the cell, and recognizing that they must share common assumptions and knowledge to do their respective jobs adequately.[4] Progress in NSB surely will serve all these and just as surely will be served by them. Moreover, the modeling aims and techniques will surely be in at least some part different, and NSB will have a lot to contribute to these areas as well as to derive from them. In large fractions of their domains, evolutionary, developmental, and genetic investigators are being forced to take a systems biology perspective, and so systems biology should grow as their methodologies spread among related disciplines. But it cannot avoid them.

3. THE PROBLEM OF DATA RELIABILITY IN THE ANALYSIS OF LARGE SYSTEMS

To illustrate how and why evolutionary and developmental concerns are central to core issues in systems biology, I want to start with a seemingly unrelated puzzle: How do we get a reliable account of the cell when we do not have totally reliable data about it?[5] This applies both to the analysis of gene-control networks and to biochemical pathways. Although the data I draw upon comes from the former context, it obviously must influence the latter. There are uncertainties

[4] Community ecology and traditional systems ecology also share many of the same methodological approaches, tools, and problems, but on more macroscopic objects. Though I do not discuss them here, they too should be a part of the broadly conceived systems biology.

[5] I was first made aware of the magnitude of the problem of data unreliability by Beckett Sterner, who also provided me with the key reference (Deane et al., 2002). Sterner's input was crucial and my debt is substantial, because this is the key organizing insight of the paper.

about the magnitude of parameter values, but even more about whether components interact at all (both under the conditions studied and in living organisms). To sketch what I will argue, fundamental features of living systems are crucial to how we can deal with these data errors, and these features require recognizing the developmental and evolutionary natures of these systems.

The new use of 'gene chips' presents an array of new possibilities in the massive volume of data produced. Unfortunately, just when we would seem to need more accurate data rather than less, they also appear to present us with much higher error rates. These DNA microarrays use miniscule amounts of different DNA sequences ('probes') to detect RNAs that may be involved in producing active proteins. Chips with tens of thousands of distinct probes are common. The latest and largest number is nearly 400 000 on a single chip. They are used not only for broad censuses of activity, but also for more targeted ones such as identifying interactions in specific metabolic pathways or disease states. The targeting is as simple as the choice of what to spot in the array.[6] They can either detect or compare activity patterns using several different protocols, but so far in a boolean ('yes/no') rather than quantitative manner.

This new technology allows an enormous reduction of labor and coordinated detection of simultaneous activity patterns involving multiple genes or proteins in a cell, something that would have been impossible two decades ago. Gene chips also increase the uniformity of assay procedures. The changes they have provided are not unlike the move from 'single-unit' recording in the neurophysiology of the late 1950s–1970s to the localization of massive changes in activity patterns by brain regions possible with functional nuclear magnetic resonance (fNMR) beginning in the 1980s. This transition produced not only new kinds of data, but also inevitably new orientations in theory and is probably responsible for the emergence of the new discipline of neuropsychology, which draws heavily on the sort of molar data produced by fNMR. The change in orientation and questions that could be addressed was enormous in both cases (though for gene chips, still early in its course), but so also in each, the increase in breadth of information was accompanied by a loss in its specific local quality. As a result, in both cases, we have the continuation of two technologies rather than the replacement of one by another.

Because the reduction in data quality with gene chips is substantial, its accuracy is a key issue. Deane et al. (2002) conducted an evaluation by comparing the results of these 'high-throughput' methods with a database of already known interacting and noninteracting proteins, producing the expression profile reliability (EPR) index. A second method (PVM) involved determining how likely the individual interactions were by asking whether they had paralogs that interacted. The second method picked up only 40% of known interactions,

[6] Wikipedia entry on 'DNA arrays', accessed 23 June 2006.

but had a false positive rate of only 1%. So it was a conservative method for inclusion of an interaction. With the two together, they estimated that of a list of 8000 protein interactions from the Database of Interacting Proteins,[7] about 50% are reliable, and using the latter test, identified 3000 of these as likely true interactions. These are chastening error rates. While-high throughput methods can generate candidates at an enormous rate, validating that they are indeed interactions requires more intensive analysis on an interaction-by-interaction basis. Moreover, Deane et al. note but do not deal with the issue of false negatives – how many interactions may be missed by the high-throughput screen.

Can we do better? One obvious move is to try to improve the quality of the data. There are many possible sources of error. We should be interested particularly in systematic ones because this fraction usually indicates problems we can do something about, by supplementing or recalibrating our methods. It may also indicate sources of systematic bias in methodological or theoretical approach (Wimsatt, 1980, 2007). Thus the high number of false positives noted by Deane likely arises at least in part for a systematic reason: Testing for the possibility of chemical interaction among all possible reactants does not allow the presumably substantial number of interactions that do not occur in vivo because the reactants are spatially or temporally segregated in the organism under natural conditions – sequestered by design. This points to the need to move beyond the current focus of NSB on intracellular dynamics. We can hope to correct these kinds of error, but only by investigating intracellular and intercellular structure and morphology, how they change through development, and how they may act to catalyze and compartmentalize reaction dynamics.

4. DATA ERRORS AND MOLAR SYSTEM PROPERTIES

Do we need to know everything? At this level of analysis, we must treat all sorts of errors as the same: reactions left out or reactions erroneously included, and assume that whether errors of either sort make a difference depends upon the context – upon the network structure of the system, which of its products are crucial, and how sensitive system behavior is to the levels of the products.

At first this suggests detailed local analyses. But at a more molar level, how sensitive system performance is to errors depends upon what kind of system we are analyzing. Software is very breakable, but consequences can be large or small. Y2K turned out to be less of a problem in part not only due to massive

[7] Some care was taken in assembling this database. About 2000 of the 8000 were identified in small-scale experiments (from 800 research articles), and the rest came from four high-throughput screens. Nonetheless, the overlap between these sources was described as ‘petite’ – the factor originally motivating the calibrations they performed.

preventative changes in software and hardware, but also because many errors were not fatal, and computer users were not unaccustomed to rebooting frozen machines already. (Taking the larger system into account thus reduced the real threat.) How about organic systems? For decades of progress in genetics, using reductionistic analyses (in which more details more accurately known were always better), we were presented with stories about how small changes in a system may change its behavior radically – and indeed that was how genetic changes were supposed to affect molar systems. The paradigm case for this view was the now classic story in every textbook of how a single base substitution in the gene coding for the beta-chain of the hemoglobin molecule could lead on the one hand (in the heterozygote) to greater resistance to malaria, and on the other to severe anemia and 'sickle-cell crisis' and an early painful death among HbS homozygotes. And that was just the beginning: There was a long list of 'single gene' genetic diseases.

This presents a picture in which organisms are like computer programs, so we need to know the genetic constitution and the biochemistry of the system in great detail because a single change could wreak major havoc. Yet software has an organization: We often do not need to know the details of a procedure to write other parts of a program, and object-oriented programming has increased the modularity of code substantially. While there are serious problems caused by inaccurate data, the picture of organic systems that would require complete knowledge to analyze any aspect of system behavior is not accurate: if it were we would be unsurvivable, unevolvable, and unstudyable. In some ways, at a very simple level we are like a house lighting system: There are things, like a failure at the main junction box, that can shut down the whole thing, and blown fuses can temporarily take out subsystems of varying sizes. But most failures in the system are, and are designed to be, strictly local. Thus individual bulbs and appliances can fail without requiring anything more than their replacement or repair because of the parallel (redundant) organization of the house wiring at the lowest level. Unreliable data are critical problems for the NSB, as they must be for the analysis of any complex system,[8] but there are various mitigating conditions.

Faced with the problem of unreliable data, we must find workarounds. Some must come through improvements in technology and the development of better means of testing the accuracy of our data. Some come through choice of questions that are less affected by this problem, though this may skew research and theory construction, as we saw earlier with the differences engendered by 'single' vs. 'multichannel' approaches, both for neurophysiology and for cell biology.

---

[8] Taylor's (1985) study in ecological communities in nature and with simulations showed that interactions left out could make studied components appear causally connected when they were not, or make them apparently independent when they were not. The conclusions would apply for networks more generally.

But significant relief can also come from the nature of the systems we study. Living systems are robust. And they are evolvable. If they were not both of these, living systems could not survive environmental fluctuations and would not have evolved in the cumulative and diversifying manner we see. Analysis of system robustness is a topic of importance in NSB (Bruggeman, 2005, Ch. 5).[9] One particularly robust (and important) feature is that organic system architectures show significant differential generative entrenchment. It is inevitable that some things have more consequences than others in the operation of a system.[10] Elements that play large roles in generating or maintaining the behavior of the system, and for which there are no alternatives, will have high evolutionary stability because their wide usefulness has rendered them irreplaceable and they are increasingly constrained in the ways or degrees to which they can change. The pervasiveness and importance of robustness and entrenchment should make them proper topics for investigation in the NSB in their own right.

They are discussed here for another reason also: They can also be used as tools for identifying and getting relevant knowledge about these systems, and in part for ameliorating effects of unreliable data. There has been an explosion of interest in the last 5–7 years in the role of robustness in the design and evolution of complex systems. This has recently been nicely reviewed and synthesized for organic systems by Andreas Wagner (2005). Robustness can also be a (usually selective) help when reliable information is hard to come by. You may not require much information about variation or exact values of variables in dimensions for which a system is robust if outcomes are relatively insensitive to the details (Levins, 1968). We can tolerate a higher rate of errors in the specification of a system in the analysis of those properties. And system analysis can tell how robust a system is and in what dimensions.

We can also get help at the other extreme: For entrenched things, by contrast, outcomes may be strongly dependent on details but that very necessity has anchored the architecture against change in those respects. Systems that do not maintain them do not survive, so these elements are relatively constant and often more readily determinable (Wimsatt, 2001). So both robustness and entrenchment are important for the analysis of system behavior, and an understanding

[9] Robustness can mean just relative stability of a system property (e.g., rate of production of a metabolic product) across different parameter values (concentrations and reaction rates) in the system, or it can mean stability across addition and subtraction of interactions or relative invariance across an ensemble of systems of a given type. To be evolutionarily stable would often require the latter and stronger conditions. Differential entrenchment, robustness, and evolvability are arguably more robust still, because they are characteristic of evolving and living systems of all kinds.

[10] To preserve deeply entrenched elements in the face of dissipative forces (mutation, etc.), it is important not only that some things be deeply entrenched, but also that some things be lightly entrenched, such that their loss is not costly. These together set up a dynamic that is crucial not only for preserving the most important elements, but also for allowing an explosion of variation that may allow exploring other optima when either internal or external conditions relax (Wimsatt and Schank, 2004).

of their implications are deeply rooted in evolutionary and developmental perspectives. I will explore the role they can have in ameliorating the problem of incomplete and unreliable data in the next two sections.

5. ROBUSTNESS AND THE MANAGEMENT OF UNCERTAINTY

The following strategies and constraints seem reasonable ways of dealing with uncertainty in data:

(1) Particularly if data is important, try to determine it in more than one way. That is, incorporate robust designs to increase reliability into your experimental methodology. This not only reduces errors through cross-checking, but can also be used to detect systematic differences that may lead to technological improvements to reduce errors and to have better knowledge about when data can and cannot be trusted (Wimsatt, 1981; Levins, 1968). The late Sylvia Culp (1995) provided powerful and revealing examples in her analyses of diverse methods in molecular genetics.
(2) Models of smaller circuits or systems require less data to keep the included errors to reasonable levels. Learn how these circuits behave, and their sensitivity and robustness to changed structure and parameter values. If they are robust, use them as 'seeds', taking their outputs as given, and investigate the circuits including and intersecting them.
(3) When modeling larger circuits, look particularly for their robust properties.
(4) Take the values and behaviors emerging from such simulations with a grain of salt. Regard the simulations as exploratory rather than definitive.
(5) After finding a behavior that is somewhat robust, try specifically to 'break' it, determining the conditions under which it fails. These might be informative: it might be a tunable switch or threshold device, breakdown conditions may indicate other variables that must be maintained or other dimensions in which it is designed to be robust.
(6) If you find a property that appears to be biologically important, and it is not robust, be suspicious of your model or the assumed parameter values. This is the complement to the old maxim of adaptive design that 'Nature does nothing in vain'.[11] The more important something is, the more important it is to guarantee its presence. Nature does not guarantee anything, but it is a good working hypothesis. So if the property is fragile in your model, explore the possibility that it is not important, that the model is wrong, or that you have misidentified its function and what it is doing.

[11] The primary application of that principle in this context is the reverse engineering one: the more complex is a mechanism, the more important is its function or functions. This may fail to be true if the mechanisms and function have been recently co-opted from another functional system, or 'kluged'.

At multiple levels, Wagner (2005) found recurrent patterns, whether it be for the conformation of an RNA or a protein, the generation of a crucial product, or the production or maintenance of a required morphology confirming this assumption of robustness: The natural system was more robust under neighboring perturbations, whether genetic, structural, or dynamic, than for values of that system drawn at random from possible systems like it. This could be for different reasons, and Wagner investigates their plausibility and scope extensively at various levels of organization.[12] One of his conclusions was that robustness to environmental fluctuations probably was the source of selection that conferred robustness to the effect of mutations as a secondary effect.[13] This is interesting: The kind of results systems biologists can get directly is relevant to evolutionary questions. More generally, one must consider that:

(1) A system property might be robust because state-space neighborhoods where a property is robust are easier to find in an evolutionary search.
(2) Once found, if the property is selectively advantageous, it is easier to maintain in the face of mutation and environmental perturbation if small induced changes in state leave the property relatively unchanged.[14]
(3) Structural changes in the system may change the character of neighboring state-spaces and, in particular, may act to increase the size of what Wagner calls the 'neutral space' in which the relevant property remains unchanged. Something like this must be going on in what Waddington (1957) called 'genetic assimilation', in which selection changes the expression of a property manifested only in the presence of an environmental shock so that it is manifested under a much wider range of conditions.
(4) The probability that something will be entrenched, that other parts of an accumulating adaptive structure should come to depend upon it, should be a monotonically increasing function of its stability and persistence. The nonlinear amplifications of selection found by Wimsatt and Schank (2004)

[12] In his discussion, Wagner focuses on the first two, though I believe that all of them come up in passing elsewhere in his discussion.

[13] My one reservation about this claim is that it turns on the fact that environmental perturbations are much more common than mutations. But if one counts recombination in a system with lots of epistatic effects as producing new mutations (as they likely will at the phenotypic level), the number of mutations goes up enormously – by orders of magnitude.

[14] Kauffman (1969, 1985, 1993) commonly assumed that the given circuits would be realized in only one specific state – that any deviations led to reduced fitness. This made them highly sensitive to mutation and was the major reason why his results seemed to establish that selection could not maintain large complex systems. The larger is the 'neutral neighborhood' for a property, the more easily it is maintained by selection. We called this 'degree of genericity' when we argued that selection and self-organization would work most effectively in concert, when the selected state was multiply realizable, and thus not radically improbable (Wimsatt, 1986, Schank and Wimsatt, 1988). Given the sizeable neutral neighborhoods for features found at many levels of organization in evolved systems (Wagner, 2005), Kauffman's claims were too pessimistic.

> suggest that this probability increases nonlinearly as well. Selection should tend to entrench robust or generic features, though one would also expect a steady accumulation of contingent, arbitrary, improbable features that become sufficiently entrenched to persist and become phylogenetically distinguishing features.

A concept used several places above, which should become very useful to NSB and is elaborated to great advantage by Wagner, is that of a 'neutral space'.[15] This is a further abstraction and generalization of the idea of 'neutral percolation surface' introduced by Huynen, Stadler, and Fontana in their important 1996 paper on evolution in RNA configuration space. Huynen et al. looked at the major forms of folded configurations of RNA molecules of length 100 nucleic acid bases and found that a few major forms dominated. They also found that the regions in which specific major forms appeared tended to be connected in mutation space, such that one could often move, one mutation at a time, throughout a connected neighborhood without changing the folded configuration, and thus, to a first approximation, remaining 'neutral', preserving the function or the fitness of the molecule. This meant that a population could 'percolate', one mutation at a time, to distant parts of the space. They also found that most major forms were reachable within a small number of mutations from one another. These 'neutral spaces' thus provided 'don't care' conditions for the composition and behavior of systems using these molecules, but the diversity of 'neutral' positions they could occupy could lead to rapid divergence if external conditions changed and selection for different configurations became advantageous. This idea itself is clearly related to the concepts of a fitness topography and to an energy surface, but is usefully exploited here as applied to discrete state systems.

Wagner's extensive discussions show that this situation or others analogous to it (as he points out, one cannot always formulate problems such that a well-organized discrete space can be defined for them) is characteristic not only of RNA space, but also at other levels of organization as well, in particular, to protein space and to spaces characterizing the dynamics of metabolic systems. This generally means that organic systems are designed so that they have many 'don't care' conditions – that behavior may often not need to be fully specified to be able to predict the dynamics with reasonable confidence. This is not

[15] Wagner's book is also distinguished by the robustness of his review of the evidence. Most claims are addressed using two or more distinct modeling strategies, often several, with the limitations of the different strategies compared. Some of the modeling concepts, for example, the idea of a 'lattice model' of a protein, in which all amino acids are of the same size and separation and their bonds can only take up angles at 90° intervals (0, 180, 270 in two dimensions, and the analogous lattice angles in three dimensions), are lovely and revealing. This one, for example, explores the consequences of topology (one-dimensional connectivity) and the distribution of hydrophobic and hydrophilic sites for the folding configurations.

something we would have a right to expect were we not talking about evolved systems, but it is bound to help with the analysis of complex systems in the presence of unreliable data.

6. GENERATIVE ENTRENCHMENT

Robustness is interesting not just because it makes organisms survivable and evolvable, but because robustness itself seems to be so pervasive among organisms – in a word, so robust.[16] This is surely why there is so much interest in studying the formal properties of networks, and also why the NSB should not see itself in contrast with evolutionary biology, particularly the new evolutionary developmental biology. Are there other features of the internal complexity of the organism which have this kind of generality? There is one that has significant implications for the behavior of networks: The architecture of development, prima facie, can be used to predict differential rates of evolutionary change in different factors, and identify constraints on how and how much they can change (though generally not the details of how they will change). This is generative entrenchment, and in particular, the differential generative entrenchment of different elements in a causal network (Wimsatt, 1986, 2001; Schank & Wimsatt, 1988, 2000; Wimsatt & Schank, 1988, 2004).

Differential entrenchment is not an accidental feature of evolutionary systems. It is generic. Nor is it avoidable in any of our engineered systems. The importance of generative entrenchment points naturally to a number of architectural and dynamical network properties, particularly redundancy and canalization (ways of getting robustness) and modularity. Each of these act to modulate and commonly to reduce its effects and magnitude. These should all qualify as general properties of interest to the NSB, but they are also of central interest to developmental genetics and evolutionary developmental biology. By collaborating in their analysis, the NSB extends its central importance to these other disciplines.

If we consider a network, a pathway, or a cascade whether of gene activity or of biochemical metabolism, different nodes are differently connected. If we draw a directed graph for the propagation of causal effects in one of these or in any mechanism – including any of the engineered products of our modern technology – we will find that different numbers of nodes are reachable from different starting points in the network. Figure 1 is a randomly constructed directed graph of 20 nodes with 20 edges, generated by computer for our first

[16] This is at least partially due to reasons suggested by Aldana and Cluzel (2003), but Wagner (2005) also provides multiple arguments to this cumulative conclusion.

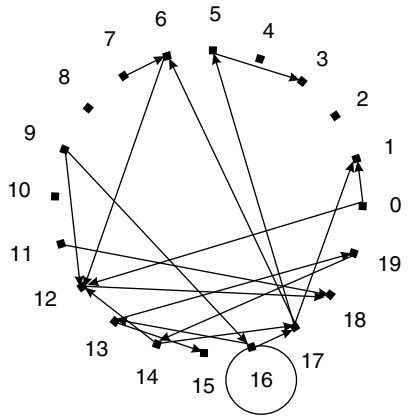

Figure 1 A directed graph representation of a gene-control network with 20 genes and 20 connections output by the program.

Nodes are genes, and a directed arrow indicates action of the gene at the tail and the expression of the gene at the head of the arrow. Mutations act on connections, and may randomly reassign the gene at the head or tail of the arrow. With 20 connections, it thus has 40 mutable sites. This particular graph was produced by 100 mutation events acting on a closed-loop model gene system of 20 genes and 20 connections. It is indistinguishable in generic properties from the one constructed at random. (From Schank and Wimsatt, 1988, p. 51. Figure copyright retained by the author.)

simulations of generative entrenchment in Stuart Kauffman's (1985) model of the evolution of gene-control networks (Schank & Wimsatt, 1988; Wimsatt & Schank, 2004). I co-opt it here to illustrate differential entrenchment.[17]

In Fig. 1 (Schank & Wimsatt, 1988) the connection from 5 to 3 has no further consequences (no arrows leave 3), but the connection from 16 to 13 has many. From node 16, we can travel: $16 \rightarrow 13 \rightarrow 19 \rightarrow 14 \rightarrow 17 \rightarrow 5 \rightarrow 3$, with other divergent paths along the way. It is rare – essentially impossible for robust reasons – to find interesting networks in which all nodes have equal influence. Differential generative entrenchment of different nodes in a network is a generic property in Kauffman's sense – virtually all networks will have it. But it is even more powerfully anchored, for it is doubly robust – generic

[17] Directed graphs provide one kind of measure of generative entrenchment, but not the only one and not always the best. Thus Wagner (2005) presents cases where volume throughput of a product at a node correlates better with evolutionary stability than the number of nodes which are reachable (downstream) in a directed graph. He also argues that the necessity of a particular product is not well represented by its topological connectivity. Many cases of the first type (where production capacity matters and is accomplished by duplicating like parts, as with multiple copies of the DNA to make ribosomal RNA, or multiple liver cells) may represent k-out-of-m structures like those discussed in Wimsatt and Schank (1988). In such cases the whole k-out-of-m structure should be treated as a single unit for evaluating entrenchment.

again under different selection regimes: Networks in which all nodes start with equal generative entrenchment will spontaneously break symmetry, generating differential entrenchment under random mutation (another consequence of genericity in Kauffman's sense, also illustrated in Fig. 1). Differential generative entrenchment will also arise spontaneously with the random addition of modifier loci or with environmental fluctuations differentially affecting different genotypes (Wimsatt & Schank, 1988; Wimsatt, 2001). But entrenchment and its conservation under selection make increasing deviations from perfect equality or symmetry (with no differential entrenchment) inevitable, and therefore self-amplifying. Loci persisting longer for any of these reasons have a greater chance of acquiring additional modifier loci, leading to their further entrenchment, and increasing disparities in entrenchment. Cellular differentiation in metazoan evolution presumably inevitably does the same thing and is crucial to the evolution of increasing size, as environmental heterogeneities for cells located in different places in the cell mass become inevitable and specialized transport and coordination mechanisms become essential.

So why should this matter? Loss of a node in a network through which many nodes are reached should cause more disruption than those leading to only a few. (The same goes for changes in its properties, which are more highly constrained if those connections are going to remain unchanged.) So, prima facie, more negative[18] selection coefficients should be assigned to changes in nodes with more nodes and connections downstream. This property is plausibly generic for causal mechanisms of all types. It may be realized differently in mechanisms of different types and appear differently in different representations of their static and dynamic structure, but it seems unavoidable.

Notice also that selective consequences and intensities emerge directly from the structural properties of the systems under consideration. This is important: when this is true, selection coefficients are not external 'add ons' to black-box models of the phenotype, as was true for population genetic models.[19] So the complaint that distances selection models from system structure is not valid for evolutionary models based in differential generative entrenchment. They are properly part of the subject matter of the NSB.

For deeply entrenched traits, the negative consequences of changing them in uncontrolled ways are virtually unconditional. The chances of making a change

[18] This simple way of putting it makes it look like a monotonic relation between variables, but actually we are talking about changes in the means and higher moments of distributions.

[19] In artificial life simulations a systematic distinction is made between simulations in which fitness measures are 'intrinsic' to the artificial organisms and those given externally, where the designer of the simulation specifies the choice rules for what is to be optimized. In the former case, the only way to tell which morph is fittest is to see if any succeeds systematically. While in principle this could be inferred from 'engineering models' of fitness, the interactions 'in practice' inevitably contain unanticipated dimensions and new unexpected consequences.

and getting away with it without changing the entrenchment are virtually zero, unless there is an array of changes that are already neutral (or nearly so) at that level of organization (Huynen et al., 1996; Wagner, 2005). To be neutral any change will usually have to meet an increasing number of constraints that generate the upper-level property that must be preserved in the face of microlevel variations in composition and process. Though the selective consequences of making changes thus are not quite intrinsic properties of the network (fitness still is a relation between system and environment), essentially no changes external to the system in the environment or selection regime will change the outcome: deep generative entrenchment is thus an 'effectively' intrinsic property of that node in the network. But as changes in the network can cause major changes in entrenchment we should be especially interested in structural or dynamical changes in the network which change the entrenchment of the network element in question. And we are interested in the connectivity patterns in networks both in general and in detail. Significant changes in generative entrenchment can emerge from addition or subtraction of a single connection that moves that part of the circuit in the direction of local integration or parcellation (Schank & Wimsatt, 2000).

How general is this? We saw that differential dependencies of components in structures – causal or inferential – are inevitable in nature. And in the symmetry-breaking, we saw that those that have some tend to get more. Their natural elaboration generates foundational relationships. New systems in which some elements play a generative or foundational role relative to others are always pivotal innovations in the history of evolution, as well as – much more recently – in the history of ideas. Mathematics, foundational theories, generative grammars, and computer programs attract attention as particularly powerful ways of organizing complex knowledge structures and systems of behavior. This is a principle of great generality, going well beyond biology to evolved systems generally. Generative systems would occur and be pivotal in any world – biological, psychological, scientific, technological, or cultural – where evolution is possible. Generative systems came to dominate in evolution as soon as they were invented for their greater replication rate, fidelity, and efficiency. We must suppose that even modest improvements in them spread like wildfire. Combinatorial generative power like that found in the genetic code, the immune system, and languages of all sorts (spoken, visual, and written) add another important dimension of amplification best treated more fully on another occasion. Information (contrary to the reductionistic talk of replicator and meme theorists) is a system property, and thus properly leads back to a properly formulated systems biology.

But no runaway processes are unbounded for which we should be thankful, else we would not be here, buried under a heap of Darwin's elephants or some

much more phylogenetically primitive sludge.[20] Generative entrenchment also cannot grow beyond limit. At some point the mutation rate (and 'genetic load') gets too great to preserve the structure, and we should expect an equilibrium between entrenchment-building and entrenchment-breaking processes. Michael Lynch et al. (1993) have analyzed this from one perspective (still within traditional population genetics) and described the behavior above the equilibrium as 'mutational meltdown'. Reduced absolute fitness from accumulating mutations decreases population size, leaving fewer possibilities to find ameliorative mutations, and the population goes extinct. His original application was to explain why asexual reproduction was not more common, but similar problems can arise for populations with mutation-sweeping sexual recombination if the overall genetic load is too great. Selection cannot maintain indefinitely large genome sizes, though various kinds of adaptations can enormously increase the size that can be maintained. (Wimsatt & Schank, (1988, 2004) consider different kinds of systemic adaptations involving generative entrenchment and genetic load that can do so.) At some point the design architecture cannot grow more: It faces a complexity catastrophe. The only escape from this is to start over again with these systems as units to build larger differentiated structures. This is the route to new levels of complexity, as Maynard Smith and Szathmár (1997) have argued, and also the route to a new hierarchical systems biology.

REFERENCES

Aldana M & Cluzel P. *A natural class of robust networks*. Proceedings of the National Academy of Sciences: 100 (15): 8710–8714, 2003.

Arthur W. *The Origin of Animal Body Plans*, New York: Cambridge University Press, 1997.

Bruggeman F. *Of Molecules and Cells: Emergent Mechanisms*. Academisch Proefschrift, Vrije Universiteit Amsterdam, 2005.

Culp S. *Objectivity in Experimental Inquiry: Breaking Data-Technique Circles*. Philosophy of Science: 62, 438–458, 1995.

Darwin C. 1859. *The Origin of Species.* Facsimile reprint, Harvard University Press, 1964.

Davidson E. *Genomic Regulatory Systems: Development and Evolution*. New York: Academic Press, 2001.

Davidson E. *The Regulatory Genome: Gene Regulatory Networks in Development and Evolution.* New York: Academic Press, 2006.

Davidson E & Erwin D. *Gene Regulatory Networks and the Evolution of Animal Body Plans*. Science: 311, 796–800, 2006.

[20] A simulation at http://www.athro.com/evo/elephs.html exhibits Darwin's illustration of the power of geometric growth with the reproduction of elephants (1859, p. 64). With a rate of growth you specify, you are invited to estimate how many years it will take until there is a sphere of elephants expanding out beyond the orbit of Pluto at the speed of light. (Not so long.) Long before that, of course, the center would have undergone gravitational collapse and sucked it all in! (Any guesses on the Schwartzchild radius for elephants?)

Deane CM, Salwinski L, Xenarios I & Eisenberg D. *Protein Interactions: two methods for assessment of the reliability of high-throughput observations*. Molecular and Cellular Proteomics: 1, 349–356, 2002.

Gilbert SF. *Ecological developmental biology: Developmental biology meets the real world*. Developmental Biology: 233, 1–12, 2001.

Huynen M, Stadler PF & Fontana W. *Smoothness within ruggedness: the role of neutrality within adaptation*. Procceedings of the National Academy of Science USA: 93, 397–401, 1996.

Kacser H. *Some Physico-chemical Aspects to Biological Organization*. Published as an appendix to Waddington, 1957, pp. 191–249, 1957.

Kauffman SA. *Metabolic Stability and Epigenesis in Randomly Constructed Genetic Networks*. Journal for Theoretical Biology: 22, 437–467, 1969.

Kauffman SA. *Self-Organization, Selective Adaptation and Its Limits: A New Pattern of Inference in Evolution and Development*. In: Evolution at a Crossroads: The New Biology and the New Philosophy of Science, edited by Depew DJ & Weber BH, MIT press, Cambridge, MA, 1985.

Kauffman SA. *The Origins of Order*. New York: Oxford U. P., 1993.

Keller EF. *Making Sense of Life: Explaining Development with Models, Metaphors, and Machines*. Cambridge: Harvard, 2002.

Levins R. *Evolution in Changing Environments*. Princeton: Princeton University Press, 1968.

Love AC. *Evolutionary Morphology, Innovation, and the Synthesis of Evolutionary and Developmental Biology*. Biology and Philosophy: 18, 309–345, 2003.

Lynch M, Bürger R, Butcher D & Gabriel W. *The Mutational Meltdown in Asexual Populations*. Journal of Heredity: 84, 339–344, 1993.

Maynard Smith J & Szathmár E. *The Major Transitions in Evolution*. New York: Oxford University Press, 1997.

Margulis L. *The Origin of Plant and Animal Cells*. American Scientist: 59, 230–235, 1971.

Morowitz HJ. *Beginnings of Cellular Life*. New Haven: Yale U. P., 1992.

Müller GB & Newman SA, eds. *Origination of Organismal Form. Beyond the Gene in Developmental and Evolutionary Biology.* Cambridge: MIT Press, 2003.

Newman SA & Müller GB. *Epigenetic mechanisms of character origination*. Journal of Experimental Zoology (Molecular and Developmental Evolution): 288, 304–317, 2000.

Oyama S, Griffiths P & Gray R, eds. *Cycles of Contingency: Developmental Systems and Evolution.* Cambridge: MIT Press, 2001.

Raff R. *The Shape of Life: Genes, Development, and the Evolution of Animal Form*. Chicago: U. Chicago Press, 1996.

Schlosser G & Wagner G, eds. *Modularity in Evolution and Development*. Chicago: U. Chicago Press, 2004.

Schank JC & Wimsatt W. *Generative Entrenchment and Evolution*. In: Fine A & Machamer PK, eds., PSA-1986: volume 2, East Lansing: The Philosophy of Science Association, pp. 33–60, 1988.

Schank JC & Wimsatt W. *Modularity and Generative Entrenchment*. In: Thinking about evolution: Historical, philosophical, and political perspectives, (eds. Singh RS, Krimbas CB, Paul DB & Beatty J) Cambridge: Cambridge University Press, 2000.

Taylor PJ. *Construction and Turnover of Multi-species Communities: a Critique of Approaches to Ecological Complexity*. Ph.D. dissertation, Department of Evolutionary and Organismal Biology, Harvard University, 1985.

Von Neumann J. *Probabilistic Logic and the Synthesis of Reliable Organisms from Unreliable Components.* In: Automata Studies, Eds. Shannon CE & McCarthy J. Princeton, N.J.: Princeton University Press: 43-98 Princeton University Press: 43–98, 1956.

Waddington CH. *The Strategy of the Genes*. London: Allen and Unwin, 1957.

Wagner A. *Robustness and Evolvability in Living Systems.* Princeton: Princeton University Press, 2005.

Wimsatt W. *Reductionistic research strategies and their biases in the units of selection controversy*. In: Nickles T (ed.) Scientific Discovery-vol. II: Case Studies. Dordrecht: Reidel. pp. 213–259, 1980.

Wimsatt W. *Developmental constraints, generative entrenchment, and the innate-acquired distinction*. In: Bechtel W (ed.) *Integrating Scientific Disciplines*. Dordrecht: Martinus-Nijhoff. pp. 185–208, 1986.

Wimsatt W. *Generative Entrenchment and the Developmental Systems Approach to Evolutionary Processes*. In: Oyama et al., 219–237, 2001.

Wimsatt W. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality.* Cambridge: Harvard University Press, 2007.

Wimsatt W & Schank J. *Two Constraints on the Evolution of Complex Adaptations and the Means for their Avoidance*. In: Nitecki M (ed.), Evolutionary Progress, Chicago: The University of Chicago Press, pp. 231–273, 1988.

Wimsatt W & Schank J. *Generative Entrenchment, Modularity and Evolvability: When Genic Selection meets the Whole Organism*. In: Schlosser & Wagner G (eds.), pp. 359–394, 2004.

Woese C. *The Universal Ancestor*. Proceedings of the National Academy of Science USA: 95, 6854–6859, 1998.

Woese C. *A New Biology for a New Century*. Microbiology and Molecular Biology Reviews: 68 (2), 173–186, 2004.