Notes on an information-processing based model of cells

Ton Rullmann, spring 1995

Introduction

Over the last seven years I have been involved from time to time in NMR studies of protein-DNA complexes. Apart from getting the structure solved, the main puzzle still is: how does selectivity arise? This can be seen as a thermodynamic problem: what are the differences in free energy of binding between various states? However, one also has to consider background (aspecific) binding, and kinetic factors. This in turn necessitates a picture on the cellular level. A priori it is not clear that specificity can be understood in static terms at all, e.g. how do you properly define the various binding states from a microscopic point of view, if their binding (free) energies differ by only 5 or 10 kcal/mol. Obviously, a microscopic MD-type calculation would have to cover an incredible amount of configuration space. Simplified interaction models might be developed to correctly predict binding preference. This would be a major feat, but would presumably shed little light on *how* selectivity arises in a real molecular system.

Setting aside the problem of describing how "bound states" emanate from a multitude of thermally accessible configurations, one could also start asking questions at a different level: how does an information processing system organize itself? Below I present a sketch of a model which is strongly inspired by Tom Ray's Tierra system[1]. Of course many people are tackling self-organizing systems, and I have no expert knowledge in the field[2] – so it may be nothing new, or unworkable, but I have to check it out. At least it was fun thinking about it.

Background

At the start I should say that the model has no straightforward connection to the recognition problem, just as there is no attempt to tie it to the molecular approach I have been using most of my professional life. The best way to explain my idea is to describe its evolution, since I am not in the position to argue for or against it.

It started in the wake of reading "Wonderful Life", where Stephen Gould argues that evolution is not characterized by linear progress, and that it does not represent a preordained development toward the emergence of the human brain and consciousness. Nevertheless, to the naïve observer it seems that at least *some* direction could be perceived, and if this direction may not be identified with progress, then maybe *complexity* is a better term. Is there a tendency for evolving systems[3] to

---

[1] http://www.his.atr.jp/~ray/tierra/

[2] When I try to trace back the origins of this note there is first of all my fascination for the problem of biological organization – which somehow should be understandable in other ways than as the outcome of just number-crunching, coupled to a long chain of "contingent details" and "weird wonders" (quoting Gould out of context): love for computers and programs, fascination for the idea of self-modifying code, some Hofstadter-Gödelian insight into the central meaning of (de)coding, the order out of chaos problem/wonder, repeated bouts of wrestling with the optimization of structures and a measured degree of dissatisfaction with MD, brief encounters with the neural net, "let there be Life to the automaton and spinglass in a protein", the magic of Alan Turing's life - who broke secret codes ánd wanted to install a white-noise generator in his first digital computer ánd showed how complex systems can give rise to complex patterns, the beauty of evolution itself - as told by Stephen Jay Gould, the incredible Tierra (and Tom Ray's personal way of writing about it), and finally this marvelous quote from Gregory Bateson: "Rather, if the world be connected, if I am at all fundamentally right in what I am saying, then *thinking in terms of stories* must be shared by all mind or minds, whether ours or those of redwood forests and sea anemones."

What it all adds up to I don't know, but I would definitely like to see a good story unfold.

[3] Initially I wanted to write "dynamical systems", but of course there is a distinction between dynamical and evolving. An evolving system is dynamical, but not every dynamical system is evolving. So what makes a system into an evolving system? The first that comes to mind is: evolution is linked to creation of order. But then of course it may

become more and more complex[4] ? Is this a general tendency, or does it manifest itself only when certain conditions are met? Is there a limit to the degree of complexity?

Trying to think of modelling an evolving system I initially imagined using generalized functions, e.g. expressing the fitness in terms of abstract parameters (just numbers) dependent on genotype and environment. But then it seemed that the model would have to specify the form of the functions too: it should specify types of relations, a physics or grammar. Specifically, the model would have to say something about gene organization and about relations between genes and environment if it were to say anything about internal complexity. Genes could not be represented as "just numbers" but would have to be assigned "meaning". Would it be possible to formulate a model which has an inherent "self-defining" capacity? [See Appendix for slightly expanded version of this paragraph].

The model

In response to these questions, and the example of Tierra, these aspects of a model came to the fore:

• Two-level system: raw molecules (resources), and organized complexes (cells). I think of both in terms of "strings", i.e. one-dimensional symbol sequences.

• Internal complexity (genes) is mirrored by external complexity (populations of interacting cells; variable resources).

• Two-way relationship: resources can be assembled into cells; cells can excrete metabolites, or decay into resources.

• Basic potential for "self-defining capacity" arises by allowing for multiple interpretations of the symbols, dependent on the context: strings may represent a resource, a coding gene, a binding site, or a regulatory element.

The basic idea would be to let the cells execute their internal program string. The program would include instructions to search for certain types of resources, metabolize them (i.e. convert them to other strings) and excrete or store them. Regulating genes would direct the program flow, e.g. by jumping and looping, and by targeting the search process.

The proposed model differs considerably from the Tierra system. The cells are more simple (no registers). The language describes symbolic transformations and associations rather than a primitive assembler code. Resources are explicitly part of the model, not just as global limiting factors. It is attempted to make no a priori distinction between "code" and "data".

The parallel to a real molecular system is clear, since there too molecules can have overlapping roles: resource, metabolite, gene, enzyme etc. It would appear though, that in real life molecules have become more specialized than in the simple model outlined above. The interesting question then of course is: how has this specialization come about?

The model appears to mix chemistry and biology into one. I don't know whether this is a complication making the whole thing unworkable, or another example of (un)specialization. The model is not an attempt to represent a real cell in terms of molecular components. Rather it is an experiment for studying organization and complexity in quasi-chemical systems.

---

equally well be related to disappearance of order. Compared to "dynamical" evolving does seem to imply an extra element of change, of (dis)appearance.

[4] Complexity could refer to internal complexity of component parts, as well as to complex composition and dynamics of the system as a whole ("diversity" and "specialization").

Topics to be investigated could include: emergence of stable patterns; complexity vs. system size and composition; extent of overlap between different functions of strings; composition of the "soup" (the set of resources): size distribution, relation to cells, specialization (e.g. small, ubiquitous vs. large "essential" strings); effects from different forms of structuring the soup, i.e. the search process.

A more complicated theme to study is the organization of strings into cells. One could start with a model without separate cells: just a long string acting on itself (probably multiple instruction pointers would reduce the chance of ending up in an infinite loop). Then in the next step some form of granularity would be introduced: restricting (or preferentially weigthing) instruction pointers and targets to be in certain areas; processing of strings in chunks; delayed output to the "soup". Conceivably certain combinations of regulating genes, directing the "program flow", could effectively act as protective cell walls.

Proceeding to a still higher level of complexity one could start to investigate interactions between cells, by allowing metabolic targets to be found in other cells, which may or may not trigger a response etc.[5] It would be interesting to see how mechanisms of communication, defense and attack might be implemented using simple regulation "genes".[6] Also more complicated regulation mechanisms could be explored, e.g. by departing from the linear program to a network representation.

Finally I note that it might be possible to use the proposed model for optimization purposes if the problem at hand can be mapped on a symbol string.

At this point I am not sure whether the model should be implemented as a GA or in the Tierra fashion, i.e. by implementing procreation via the instructions themselves instead of an external fitness function and selection scheme.[7] In both cases it would be make sense to make procreation dependent on the cell's ability to reproduce its own string(s) by metabolizing resources.

Numerous practical issues have to be worked out: the search procedure (sequence, cost); gene reshuffling and mutation; errors (mismatches/unability to find target, metabolic waste); the metabolism implementation (time allocation, internal storage, fragment combination, energy function?), and of course the coding.[8]

---

[5] In Tierra boundary crossing is an essential feature, but since there is no metabolic processing going on, it is code rather than resources that are being shared between cells, giving rise to interesting parasitic and social models of behaviour.

[6] Writing this I wonder what the proposed model could have in common with cellular automata worlds, where self-replicating structures can exist. Both systems use rule sets and coded states, but in CA the states are assigned to spatial points, while in a Tierra-like system they are assigned to logical entities.

[7] Maybe the difference between these two approaches is not really fundamental. In a Tierra-like model one still has to add a schedule for removing cells to deal with overall memory limitations. Such a scheme – whether based on error accumulation, production of key-compounds or anything else – effectively defines fitness. Similarly, one needs an algorithm to allocate processor time, and this influences competition and therefore survival.

[8] Symbols in resource (raw data) strings, target (address) strings and code (instruction) strings could all occupy the same symbol space. Targets could be specified as identical to the actual resource sought, or as its "inverse", or another variant (wildcards?). Different target lengths may be embodied in the instruction set (instruction followed by target; optional 'end' symbol). Some instructions could direct the search process (initial location, length). Other instructions deal with metabolism (specify point of action; excise, insert or mutate symbol(s); store, combine, or excrete fragments) and with 'gene coordination' (move, reset or change direction of instruction pointer). All instructions could be represented by different symbols or pairs of symbols, or several instructions could be mapped on the same symbol (using separate 'signal codons' to determine the context).

Gene-organization, or 'program structure', is a separate issue. The simplest, and most Tierra-like approach, is to have one instruction pointer moving incrementally from symbol to symbol. Another, more chemical approach would be to start modeling the 'program' as a network of coupled pathways (multiple pointers? program granularity? kinetic matrix with varying rate constants?).

Summarizing, I can see three different themes to be studied:

a) evolution of complexity

b) evolution of functionality (overloading or separation of functions, choice of functions, relation to organization of ‘programs’ and ‘resources’)

c) evolution of organization (transition from strings to cells; interaction between cells; internal program as network)

Appendix

Stephen Gould argues in "Wonderful Life" that evolution is not characterized by linear progress, and that it does not represent a preordained development toward the emergence of the human brain and consciousness. Instead of a "cone of increasing diversity" characterized by stately progression to "more and better", Gould sees initial disparity followed by decimation and specialization. Evolution is not a process with predictable long-term results but is shaped by, and as, a long chain of contingent details.

Nevertheless, to the naïve observer it appears that at least SOME direction could be perceived, and if this direction may not be identified with progress, then maybe *complexity* is a better term. Is there a tendency for evolving systems[9] to become more and more complex[10] ? Is this a general tendency, or does it manifest itself only when certain conditions are met? Is there a limit to the degree of complexity?

When I try to think of modelling an evolving system in biological terms, concepts such as genotype, fenotype, mutation, fitness, species and population come into play. It would e.g. be possible to define a function expressing the fitness in terms of parameters dependent on fenotype and environment. The fenotype would be a function of genotype and environment (and possibly time) etc. While the *values* of these functions could be varied within arbitrary bounds, the crucial thing is that the model would have to specify the *form* of the functions: somehow the model has to specify *types* of relations, a physics or grammar if you like. Specifically, the model would have to say something about gene organization and about relations between genes and environment if it were to say anything about internal complexity. Genes could not be represented as "just numbers" but would have to be assigned "meaning".

This is tantalizing, since in my view evolution is inextricably linked to appearance of meaning: nucleotides becoming members of the genetic code; molecules acquiring a role in a reaction network; DNA segments becoming regulator elements; an inconspicuous novel trait turning into a species-saving property. There is a strong aspect of self-definition here: components acquire their role within a constantly evolving system which is being shaped by the components themselves. Concepts such as complexity, diversity, fitness and information all refer to a context, or a classification of contexts. If the contexts change, so does the information.

So this is the fascinating aspect for me:

Would it be possible to formulate a model which has an inherent "self-defining" capacity?

(not implicit, through hierarchical buildup, but as explicit possibility)

[9] Initially I wanted to write "dynamical systems", but of course there is a distinction between dynamical and evolving. An evolving system is dynamical, but not every dynamical system is evolving. So what makes a system into an evolving system? The first that comes to mind is: evolution is linked to creation of order. But then of course it may equally well be related to disappearance of order. Compared to "dynamical" evolving does seem to imply an extra element of change, of (dis)appearance.

[10] Complexity could refer to internal complexity of component parts, as well as to complex composition and dynamics of the system as a whole ("diversity" and "specialization" in biological terms).