

Questions of function: modelling the emergence of immune response.

Tom Hebbbron and Jason Noble

University of Southampton, Southampton, SO17 1DJ, UK,
tomhebbbron@zepler.net

1 Part One: Theoretical Background and Motivation

The practice of borrowing designs from nature to solve problems is increasingly prevalent in our technology development. The reductionist paradigm that allowed technology to advance to its present state fails when faced with many richly interacting components: a hallmark of complex systems. The behaviour of such systems is not solely determined by the aggregate behaviour of its components, but by emergent effects. The divide-and-conquer strategy is insufficient as the whole is greater than the sum of the parts.

Happily, the natural systems that surround us have been solving this class of complex problem for several billion years, optimising solutions through natural selection. By idealising these solutions we can aim to transfer successful strategies from biology to our own problem domains.

When borrowing ideas from biology in this way, there can be pitfalls in both under- and over-reaching. For example, an artificial immune system (AIS) algorithm that performs binary classification (e.g., "safe" / "dangerous") is not really doing what a real immune system does. The algorithm may do well at the practical task we set for it, but it is not embedded in real biochemistry in the way an immune system is. There is a superficial similarity but we would be wrong to draw conclusions from this understated algorithm about real immunobiology. Conversely, building a high-definition model of immune system components might be a valuable scientific goal, but as engineering it would be a misplaced effort if the complex contextual constraints of the real immune system were not present in the target problem. It is important to look closely at the differences between the practical problems tackled by AIS researchers (e.g., computer security) and the biological problem faced by real immune systems, as a false assumption of congruence will lead to either models that are too simple to be good science, or tools that are too byzantine to be good engineering.

This is the real crux of the problem: how to separate the useful functionality of a biological system from the legacy of its evolutionary history. Whilst evolution may appear to produce general purpose solutions, this is not the case; solutions in real biological systems are *nichiversal* not universal. Organisms evolve to solve the problems their ancestors faced, but not to solve a general class of problem [1]. Solutions are tightly bound to the context in which they arose, and represent a compromise balancing multiple conflicting constraints on the

organism. A fundamental constraint is that the solution is made of the same biochemical substrate as the rest of the ecosystem it exists in. This is not the case in the relationship between artificial immune systems and the material that they work on; the AIS is able to step outside of the constraints of the data it is working with, breaking out of the bounds of the reflexive relationship that underlies biological immune systems.

This question of compatibility between a biological solution and a computation problem is a question of function, one of Tinbergen's "Four Questions": a schema of four equally valid categories of enquiry into biological or ecological systems [2]. In brief:

- **Mechanism:** how does the proximate system work? What are the component parts, and how do they interact?
- **Ontogeny:** in an organism, how does development of the mechanism proceed from origin to maturity?
- **Phylogeny:** what is the evolutionary history of the mechanism and its ancestors? What innovations appeared, what scaffolding disappeared? What is the particular path taken through the space of possible phenotypes by the ancestors of the modern organism?
- **Function:** what is the selective advantage in being equipped with such a mechanism? What ecological problem does it solve? The adaptations in the phylogenetic route taken by a species are specific responses to selection pressures. A truly explanatory account of an evolved mechanism must address the deeper structure of the problem for which the specific mechanism is one possible solution.

Questions of mechanism and ontogeny are of primary interest to the medical sciences, and it seems a great deal of AIS work draws from these models of proximate mechanism. Comparative genomics helps answer questions on phylogeny: when and where did recombination activating genes (RAG) arrive? How conserved are pattern recognition receptors (PRRs) across the eukaryota? The question of function however, is often presumed to be implicit — immune systems are *for* protecting against pathogens. This question warrants further exploration: for there to be pressure for an immune system in the first place requires certain dynamics in the ecology to hold — pathogenic behaviour must exist, and defences against it must be available and evolutionarily findable. Once an immune system emerges it makes fundamental changes to the future evolutionary pathways available to a lineage; some adjacent possible phenotypes will be incompatible with the defenses now encoded in the genome.

Why should we be interested in answering questions of function? High-resolution, predictive models of biological mechanisms are certainly useful, for instance in developing therapeutics; but such models are opaque. We can see how they work, but not why they work the way that they do, and such models are open to misinterpretation. For instance, low iron levels or a fever may not be a symptom of infection but an immune response, creating an unfavourable environment for a pathogen. Without knowing the function of this mechanism, there is the danger that therapeutics are employed to lower the temperature or

restore iron levels, working against the biological mechanisms [3]. This need for functional explanations to understand a mechanism in context is also an example of the "no free lunch" theorem: in short, you cannot evaluate an algorithm for solving a problem without some knowledge of the problem and how it arose, and there are no universal solutions — the success of an algorithm on one problem does not necessarily translate to another. This is significant for artificial immune system practitioners, who are trying to isolate the design principles of a system particularly tightly embedded in its biological context.

2 Part Two: Details of the Proposed Model

To investigate functional questions on the evolution of the immune system, we must consider the major transitions in immunity, and why they arose. We see these transitions as follows:

1. The protection of metabolism from random external perturbations (the vast majority of which are fatal), through some form of membrane.
2. Predation as a strategy; exploiting other organisms as local concentrations of resources and the corresponding coevolution of defensive counter strategies.
3. The accumulation in the genome of these defences — the emergence of innate immunity. The continuing selective advantage of these defences relies on the conservation of features of pathogens; features crucial to their metabolism or membrane and difficult to mutate to non-recognised alternatives.
4. Asymmetry of pathogen / host evolutionary rate, and an increase in pathogenic load selecting for the emergence of somatic mutation: adaptive immunity to counter within-lifetime diversity of pathogens.

The emergence of a membrane to isolate an auto-catalytic set from the random perturbations of the outside world will be assumed in our model, as modelling at the level of physicochemical interactions that would be required to capture the emergence of membranes and proto-cells is below the level of abstraction that we believe will be necessary to observe higher level strategic events.

The substrate that our ecology will be built from is metaphorically closest to the metabolome in real biology. We require a level of abstraction that is higher than that involved in building an artificial chemistry, but that still allows for an organism to be made up of genetically specified phenotypic components that interact to provide energy payoffs. The model should be general enough to represent differing levels of epistasis in the landscape of energy payoffs for these component interactions.

We took inspiration from Kauffman's NK model of tunably rugged fitness landscapes [4]. An N by N interaction matrix defines the energy gain or loss $[-1.0, 1.0]$ of all possible interactions between metabolic components. There is no concept of concentration, so a metabolic 'state' in the model is represented simply by a bitstring of length N , which serves as an index to the interaction matrix: a 1 or 0 indicates the presence or absence of a particular metabolic complex.

Fitness will depend on the ability to reproduce, a function of genome size and energy gained from interactions. A membrane will be given in the model, and organisms will be represented by two binary indices to the interaction matrix: internal and surface. The internal components are able to interact with the environment only via the surface. An abiotic background solution of nutrients will also be represented by a binary index to the interaction matrix. Energy gain or loss at each timestep is a function of the interaction between the internal and surface components and the surface and external state. Death occurs when energy falls below zero, and reproduction occurs when energy is sufficient to divide, and have both mother and daughter remain alive.

Reproduction with a small mutation rate will see adaptation; the population optimising surface and internal components according to the interactions available from environmental nutrients: a search for optima on a fitness landscape whose ruggedness is defined by the distribution of values in the interaction matrix. Of course, organisms do not exist independently, and Kauffman extended the NK model to consider coevolution in the NKC variant, where fitness landscapes are dynamic; deformed by the exploration of the linked fitness landscapes of other species. This separation of the fitness landscapes breaks with our idea that the substrate must be common to all members of the ecology: we want to see life converge on some parts of the landscape, and become contingently locked into others.

We expect that in some regions of parameter space, a number of the components available will be universal; they have such a high payoff that individuals employing them will dominate the population, and find it hard to mutate away from this dependence. The interesting dynamics will come from secondary components, those that have a high payoff, but are not universal. When contingencies of evolution lock these components into a species genome, and in particular if they are used as surface proteins, they become potential pathogen associated molecular patterns (PAMPs): conserved features that are difficult to mutate to unrecognised alternatives because of their importance to the pathogen's metabolic network, but not present in non-pathogen surfaces.

Parasitic and subsequent coevolutionary behaviour has been observed in artificial life systems before, notably in the Tierra system [5], which was not explicitly designed to exhibit the phenomenon. Parasitism in Tierra is an interesting example for us because it shows that a fixed and quite brittle substrate defined by Tierra's virtual machine instruction set, and which represents a single point in our parameter space, gave rise to parasitism and the beginnings of an immune response. This gives us confidence that regions of biological interest do exist in our own modelling space and will be found without too much difficulty.

We have several different goals in running the simulation. First we would like to find regions in parameter space that give rise to the evolution of a coevolutionary arms race between parasitic and host species, i.e., the beginnings of an immune response.

Our procedure will be to initially run the ecology with no organism / organism interactions — they optimise to background nutrients, essentially finding

optima on a static fitness landscape. We will then switch on organism / organism interactions, allowing the emergence of predatory behaviour. Under what parameters do we see some species become predatory - i.e. now specialise to feed on other species rather than the background nutrients?

Do we see a coevolutionary arms race between species? Or is it transitory? What regions of the PPI parameter space lead to the behaviours of interest?

We expect to see that species well adapted to the background nutrients end up on different optima than those who experience the pressures of predation and immune development. If we begin with a static landscape and allow species to explore and find optima, they will end up on different peaks when the landscape remains static (no interactions) than if the landscape has undergone a dynamic phase (the introduction and later removal of interactions and immune response). This will illustrate the way in which the pressure to adopt an immune system drives species along different phylogenetic routes and ultimately to different optima.

In future work, we intend to extend this modelling methodology to explore the functional questions that underlie the fourth transition identified; from innate germline immune responses to somatic adaptive responses.

References

1. Bullock, S.: The fallacy of general purpose bio-inspired computing. In Rocha, L.M., Yaeger, L.S., Bedau, M.A., Floreano, D., Goldstone, R.L., Vespignani, A., eds.: Tenth International Conference on Artificial Life, MIT Press, Cambridge, MA (2006) 540–545
2. Tinbergen, N.: On aims and methods of ethology. *Zeitschrift für Tierpsychologie* **20** (1963) 410–433
3. Nesse, R.M., Williams, G.C.: *Why We Get Sick: The New Science of Darwinian Medicine*. 1 edn. Vintage (January 1996)
4. Kauffman: *The Origins of Order: Self-organization and Selection in Evolution*. Oxford University Press (1993)
5. Ray, T.S., Xu, C.: Measures of evolvability in Tierra. *Artificial Life and Robotics* **5**(4) (2001) 211–214