Chapter ?

Success and Failure in Adaptation

Anne-Marie Grisogono
Defence Science and Technology Organisation
Australia
Anne-Marie.Grisogono@defence.gov.au

Adaptation is a powerful mechanism displayed in many forms by living systems. It underpins evolution of species, animal learning, the development of culture, the response of the immune system and human problemsolving to name a few. It can also fail in many ways, as in species extinction, the development of phobias, or premature convergence on poor solutions. In order to effectively exploit the principles of adaptation to better handle complex decisions and problems we need to understand the conditions and factors for success and failure of adaptation mechanisms. We will discuss some observations and insights about what those conditions and factors are in the context of our conceptual framework for adaptation, and propose some research directions to further develop our understanding.

1 How Does a Living System Succeed in Making a Living?

A living system is an open system – it needs to import resources from its environment to function and to reproduce, and it fights a constant battle against entropy and decay by active processes of homeostasis, self-repair and exporting waste. Making a living is generally hard, but how hard it is depends on both its own properties – what the system's needs are and what it is able to do, and on relevant aspects of its environment's properties – specifically, how difficult is it for the system in this particular environment to find and use sufficient resources, while avoiding or protecting itself from lurking dangers? What problems does it need to solve to survive and reproduce?

In the real world of living systems, we see many examples illustrating how complexity and adaptation are intimately and inextricably linked, one breeding the other in mutually reinforcing cycles that simultaneously increase both the challenges that living systems face, and their capacity to deal successfully with those challenges. The former is a measure of the complexity [1] of a system's environment, the latter is a measure of its ability to adapt.

This suggests something rather profound – that the linkage between the two is no mere accident, rather that it is precisely the problems that the environment poses for the system's survival that stimulate the process of adaptation and elicit the emergence of new structures and behaviours to solve those problems. These in turn lead to increased complexity for the other living systems that are part of, and share, its environment, whether through decreasing their access to, or increasing competition for, resources, or through posing more direct threats to them, and they in turn will seek through their own adaptation processes to regain lost ground, or improve their share of resources. And so begins another round of co-adaptation … obviously in this kind of scenario, there is no 'final solution', no guaranteed enduring dominance of one system over its physical and ecological environment.

However we also observe that adaptation does not even guarantee transient success. There are many examples [2,3] of partial or catastrophic failure in living systems' attempts to adapt to changes in the environment, to increased competition for resources, or to the stresses of day-to-day problems. The results of such failures of adaptation appear in every domain, from species extinction in evolution, to misdirected learning resulting in phobias and superstitions, from premature convergence on poor solutions resulting in disappointing or disastrous outcomes, to history repeating itself because lessons observed are not learned, from autoimmune diseases and our unfortunate predilection for foods high in fats and sugars, to stock market crashes and overexploitation of resources.

We ourselves are living systems, as are our social groups, our organisations and the complex enterprises we create [4]. The rich spectrum of interactions between living systems in the struggle to survive and thrive – from predation, competition, and parasitism, through co-existence, and to symbiosis and cooperation – is as much part of our landscape as it is of any other section of the natural world, and the ability to adapt is inherent in our nature, and in the larger systems we create. Therefore the above observations also apply to us – there may be pauses, but there is no end to the struggle to adapt except through total failure, success is never certain, there are many ways to fail, and at the same time the need to be adaptive is constantly increasing as the complexity of our own environment continues to increase – largely through our own doings.

It is important therefore for our own survival and welfare, to really understand the principles of adaptation well enough to tip the scales towards more success and less failure.

2 A Conceptual Framework for Adaptation

As a first step to that end, we have developed a conceptual framework for adaptation [5,6] through a systematic analysis of how it naturally arises and operates in living systems, and have identified some of its features and mechanisms. This framework is based on the most basic concept of a living system as having the capacity to sense and act conditionally on its environment as a result of what is sensed, and a generic model of adaptation, which comprises the four essential components of:

- a concept of ‘fitness’ or relative success and failure,
- a source of variation in some parts of the system,
- a means of testing the variations produced for their impact on fitness, and
- a selection process which preferentially retains variations which enhance fitness and discards those decrease it, resulting in
- the encoding of information into the system, in a way that tends to increase the relative success of the system .

Examining the various ways in which such a generic model can be instantiated has led us to a number of classifications for specific examples of adaptive mechanisms:

- whether it operates on individual systems (as in individual learning) or on populations (as in evolution of species),
- what parts of the system are subject to adaptation – leading to five nested levels of adaptivity, which we describe as:
  - Level 1. Action-in-the-world – tuning existing sensing, processing and action capabilities;
  - Level 2. Learning – expanding or modifying existing capabilities;
  - Level 3. Learning-to-learn – improving the effectiveness of learning;
  - Level 4. Defining-Success – improving the alignment of selection with real fitness, and
  - Level 5. Co-Adaptation – with two parallel threads:
    - for those systems within our own control – tuning the allocation of roles and resources to them in a System-of-Systems context, and
    - for systems we interact with but do not control, choosing our options with a better appreciation of their consequences through anticipating their likely adaptive responses to our actions.
- what kind of change or stress is the adaptive mechanism targetting – leading to four classes of adaptivity which we describe as:
  - *Responsiveness* – ability to recognise and deal effectively with immediate threats and opportunies as they arise,
  - *Agility* – ability to recognise when changes in the context or in system capability require a change of strategy, and to implement it easily,
  - *Flexibility* – ability to recognise and deal effectively with new challenges, i.e. to be reconfigured in different ways to do different things, under different sets of conditions, and
  - *Resilience* – ability to recognise and deal effectively with harmful changes to the system itself i.e. to recover from or adjust to misfortune/damage, and to degrade gracefully under attack or as a result of partial failure, and finally,
- the scale [7] at which it operates.

These classifications offer a rich matrix of templates for possible adaptive mechanisms which we can use to either recognise, analyse and tune existing mechanisms, or to identify and exploit opportunities for engendering new ones.

To complement these classifications, and assist in creating effective interventions to improve the success of existing adaptive mechanisms, or in the even more difficult task of engendering new ones, we also need to understand how to assess and tune:

- the integrity and ‘health’ of adaptive mechanisms,
- their temporal dynamics in relation to relevant processes in the environment, and
- whatever other factors will influence their degree of success or failure.

The process of developing this conceptual framework, and of testing and refining it through applying it to various complex problem domains, has already generated many important insights about how we can exploit adaptation. We have learnt for example that adaptation processes that work on populations and those that work on individuals have complementary strengths and weaknesses – the former can explore large parameter spaces to produce new design features in the population, but are slow and wasteful, while the latter may be much faster but are best at tuning a small number of design parameters within the individual system to produce more effective use of the system’s existing capabilities.

A particularly significant insight has been the realisation of the critical role that the concept of ‘fitness’ plays in adaptation, in defining an axis for the selection process and thereby steering the system through repeated cycles of adaptation towards some parts of the space of possible outcomes and away from others. What the ‘fitness’ in fact is for a given adaptive system can be deduced by observing how the selection process actually operates. This may not be the same as what an observer might consider to be the ‘ideal’ fitness of the system, so we need to distinguish between the fitness implicit in the selection process, and an observer’s concept of the fitness of the system in relation to its intended or perceived roles and functions.

This brief overview of a conceptual framework for adaptation also makes apparent that there are a large number of parameters that characterise any particular instance of adaptation, just as there are for the system itself. As we will see in the following sections, the fitness or success of an instance of adaptation will depend not only on its own characteristic properties but also on the properties of the system and of the context it acts in as well. What we will now begin to explore is what success and failure mean in the context of adaptation, and how we can shift their balance in our own favour.

3 Measures of Success and Failure

The success of a system is so intimately bound up with the success of its adaptation mechanisms that it is difficult to disentangle the two. We can construct hypothetical cases where a system’s unearned success is due to pure chance (the environment happened to change in a way that made it easier for the system to make a living) or conversely a system is destroyed by a catastrophe (such as a meteorite impact or earthquake) that no adaptation could possibly have averted. But the rest of the time, it is adaptation that can take credit, or has to shoulder the blame for the system’s fortunes.

Even when the environment is stable and there is no great pressure on the system, we have to acknowledge that although it may be ticking over very nicely without needing to change at that particular time, it was through earlier adaptation that it acquired the capability to thrive in that environment, and that when the period of stability comes to an end, its survival will once again depend on its adaptive properties.

Moreover, adaptation is not only about innovation, it is also the mainstay of homeostasis – keeping critical parameters within effective operating ranges, protecting useful properties and structure from decay through harmful variations creeping in – and to that end is as essential to a system's success in periods of stability as during change.

So if the success of adaptation is to be judged by the success of the system it operates in, then we should start with defining success and failure for a system.

In fact it turns out that total system failure is much easier to define than total success – a system that irretrievably stops functioning or dies, or a species that goes extinct leaving no descendants has clearly and unambiguously failed. Given that our current understanding of cosmology suggests that the universe is likely to either end in a Big Crunch, or peter out into a "heat death" [8], this is the final fate that awaits every system. So one possible measure of success of a system is the time it takes to fail, but this measure is not very satisfying since it cannot be determined until the system fails, and for a system that is still functioning, we can only say that it has not failed yet.

On the other hand, there are obviously varying degrees of success, even if they are all ultimately doomed. We will now delve into these, firstly for populations of systems, where the adaptation process is evolution, and then for individual systems, where adaptation takes the form of learning.

3.1 Success of a Population of Systems

In the case of the evolutionary adaptation of a population of living systems, fitness is unambiguously equivalent to reproductive success – which can easily be assessed in many ways. The following are all snapshot measures at a point of time:

- MoS(P) 1. the size of the population,
- MoS(P) 2. the rate of growth (or stability) of the population,
- MoS(P) 3. the fraction of environmental resources that the population uses, or
- MoS(P) 4. the fraction of the physical environmental space that it occupies, or more ambitiously
- MoS(P) 5. the fraction of the habitable fitness landscape that it succeeds in finding and populating.

The first two are straightforward absolute measures, but they do not convey how much room there might be for improvement, so they are not as informative as the relative measures which follow them and score the system's performance against some kind of theoretical limit, be it simply resource-based, or more abstractly, in terms of the in-principle exploitable niches that might exist in the ecosystem.

In this latter concept we are building on the notion of reproductive success as the systems' fitness, arguably reflected in their 'ease of making a living', which as discussed above, depends on the properties of both the system itself, and of the environment.

Fitness can be conceptualised as a hypothetical surface in design space [9], the space spanned by the system's design parameters, with the 'vertical' dimension being the 'fitness', and the actual form of the surface depending on the relevant environmental parameters.

For ease of visualisation, fitness landscapes are often portrayed as surfaces in 3-D, in other words as single-valued functions of only two system parameters, but for real complex systems the number of design parameters is enormous, creating a vast hyperspace of possibilities, so this picture is of course a huge oversimplification [10], compounded by the misleading ability to ‘see’ the entire surface, and by the surface appearing to be static, whereas in fact it is only the immediate neighbourhood of currently occupied regions of design space that can be ‘seen’.

Moreover, because of impacts on the physical environment and co-adaptation of the other systems in the ecosystem, the shape of the fitness surface is highly dynamic – it changes as soon as the population of systems starts to ‘move’ in the design space, or environmental parameters change for any other reasons apart from the behaviour of our population of interest. So far from the cosy image of populating the fertile slopes of fitness peaks and peacefully spreading through the landscape to nearby peaks, a more fitting image might be of a hapless swarm endlessly pursuing and precariously balancing on the constantly shifting high ground.

Since fitness is the ‘vertical’ dimension, one can think of (hyper)contours in this space demarcating regions of different levels of potential fitness, and a threshold of habitability below which survival is not viable. A population of systems might occupy several regions of adequate fitness, but if they are far enough apart in design space they might be better thought of as distinct subpopulations or even different species if they cannot interbreed. However there may also be regions of adequate fitness which are not populated and are therefore potential niches in the ecosystem that could be colonised, if there was a way of getting there.

So, for all its weaknesses, the fitness landscape concept is a useful construct because it does open up the notion of populating a region of design space, the extent of the region representing the spread of values of the different system parameters, and the process of evolution then being seen as migration of the population in design space towards regions of higher fitness. Over time, populations carve out trajectories in design space, subpopulations may become separate, diverge, die out or spawn new branches and so on – the familiar evolutionary branching and speciation process.

The utility of this abstraction is in making apparent the influence of the topology of the fitness surface on the accessibility of unpopulated niches, and in thinking about what may impede or facilitate how a population moves in this space. It also suggests additional success measures that capture whole-of-life aspects of populations:

MoS(P) 6. looking forwards from one point in time for a given population – the number and integrated biomass of subsequent populations spawned, or alternatively measures of the diversity and extent (absolute or as fractions of what is available) of regions of the ecospace successfully populated;

MoS(P) 7. looking backwards from one point in time for a given population – its cumulative ‘productivity’ so far, for example in terms of the production of useful new features.

Another important perspective on success of both the population and its adaptation processes concerns its resilience with respect to hypothetical changes in the context:

MoS(P) 8. robustness of the population to various pertubations in its environment, expressed at tolerances on relevant environmental parameters.

What becomes important in this view are two things:

- the ability of a population to protect and expand the regions of design space already populated, through
  - stabilisation of useful properties against harmful variation,
  - being tolerant to external changes, and
  - acting on the environment in a way that increases its habitability, and
- the ability of the population to move in design space, away from regions of reduced or reducing fitness and towards regions of greater or growing fitness.

This richer picture of evolving populations suggests measures of success for the adaptation processes underpinning the success of the population they act on:

MoS(A) 1. the speed with which the population can 'move' in the fitness landscape to get off regions of declining fitness and find new suitable regions as they emerge;

MoS(A) 2. the ability to stabilise and protect useful properties;

MoS(A) 3. the ability to modify the environment to increase its habitability locally to the current population, and to foster the emergence of habitable regions elsewhere;

and corresponding measures of failure for adaptation:

MoF(A) 1. loss of agility in design space eg by overspecialisation or reduced diversity;

MoF(A) 2. loss of useful information – leading to 'mutational meltdown' [11]

MoF(A) 3. the propensity to act on the environment in such a way as to decrease its habitability either locally, or elsewhere.

3.2 Success of Individual Systems

When it comes to how the success or failure of individual systems are judged, the picture get a little muddier because there is more scope for divergences between the individual's 'ideal' fitness (which is open to many interpretations), the fitness which drove the evolution of the learning mechanism (which is not), and the fitness that is implicit in the resulting way that selection operates in learning (which is in between).

There are several reasons for this. Firstly since the real impact on reproductive fitness may not materialise quickly enough to drive learning, proxies for it are evolved to do the job. Secondly, since the individual system is not the unit of selection in the evolution of the proxies, the result of selection need not perfectly reflect the individual's interests, and thirdly in a rapidly changing environment the proxies will not be able to evolve fast enough to keep pace, and so divergence between proxies and individual ideal fitness can increase significantly. As in the case of populations of systems, an individual system can at best only achieve transient success relative to others, and will constantly strive through adaptation to increase its share.

With these caveats in mind, we posit the following measures for the success of an individual system,:

MoS(I) 1. *present* state of 'health' of the system – its integrity, level of functioning, how much above the survival line it is, and hence the amount of 'free' resources which it could use to shape the environment to its advantage;

MoS(I) 2. *history* cumulative products of the system – offspring, or 'useful' output in relation to its purpose for designed systems, MoS(I) 3. *future* option space available to the system – eg trust and respect others accord it, resources controlled, acquired capabilities and strategies;

and since co-adaptation processes and environmental change create a dynamic context, its success in these measures will also depend on the ability to adapt i.e.learn, and thereby modify the above over time. So we arrive at similar measures of success and failure for its adaption mechanisms as we did for those of populations of systems:

MoS(A) 4. the speed with which the individual can 'move' in its own fitness landscape and replace capabilities of lower or declining fitness with new ones better suited to creating success in its context;

MoS(A) 5. the ability to stabilise and protect useful properties;

MoS(A) 6. the ability to modify its environment to maintain or increase its local habitability, or to foster the emergence of habitable regions elsewhere;

with corresponding measures of failure for adaptation:

MoF(A) 4. loss of agility eg by overspecialisation or getting set in its ways;

MoF(A) 5. loss of useful information – forgetting what it has learnt;

MoF(A) 6. acting in ways that decrease habitability either locally, or elsewhere, eg by prioritising short-term gains over longer-term consequences.

4 Factors that Drive Success or Failure in Adaptation

We have seen that for both individual systems, and populations of them, adaptation is about protecting and increasing their fitness, and that since fitness is a function of both system properties and context properties, it can do this by either acting on the system to change some of its properties – i.e. move it in design space, or by acting on the context to change the height of the fitness surface in the region of design space that the system already occupies.

The effectiveness of adaptation then depends on the factors that determine the rate of exploration of design space: the rate at which variations are produced, the scope, depth and concurrency of variations, the rate at which the variations can be evaluated for their impact on fitness, the accuracy of the evaluation, the strength of the selection pressure, and the accuracy and tolerance of the selection process, and on the effectiveness of the exploration: what can be varied, whether the variations are random or targetted, and what the mapping is between the things that can be varied, and the resulting changes in system properties – the so-called genotype-phenotype map [12] which depends on such properties as the modularity of the map, as well as topological properties of the fitness landscape that is being explored [see eg 13] which will determine what parts of the fitness space are accessible, and dynamical and other properties of the context which determine the rate and types of challenges to be met.

We now see that the number of parameters defining an adaptive mechanism is even greater than section 2 suggested – essentially constituting an 'adaptation design space' comparable to the system design space we discussed in section 3.1. The task of choosing those parameters to achieve successful adaptation, is thus a complex design task in itself, analogous to the task that evolutionary adaptation performs for populations of systems. We conclude that this calls for a similarly adaptive approach to engendering successful adaptation – which reinforces the need for multilevel adaptation as described in section 2.

References

[1] This is not an absolute measure of complexity of the environment, but rather its complexity as ‘perceived’ by the system in terms of the precision required of its sensing, discrimination, and response, and the difficulty of the problems it has to solve, where complexity of a problem is given by the ratio of the number of ways of getting it wrong to he number of ways of getting it right.

[2] DÖRNER, DIETRICH *The Logic of Failure*, Perseus 1996 discusses failure due to human error

[3] DIAMOND, JARED *Collapse – How Societies Choose to Fail or Survive*, Allen Lane 2005

[4] MILLER, JAMES GRIER *Living Systems*, University of Colorado, Denver, 1995

[5] GRISOGONO, AM *Co-Adaptation*, invited paper 6039-1 Complex Systems Conference, SPIE Symposium on Microelectronics, MEMS and Nanotechnology, Brisbane December 2005

[6] GRISOGONO, AM *The Implications of Complex Adaptive Systems for Command and Control* CCRTS, San Diego, 2006

[7] There are several somewhat related aspects – hierarchical scale, scale of what is sensed and responsed to, timescales of the adaptive cycle and so on. See for example Bar-Yam, Y., *Multiscale complexity / entropy*, Advances in Complex Systems 7 (2004), 47–63.

[8] http://en.wikipedia.org/wiki/Ultimate_fate_of_the_Universe

[9] DENNETT, DAN *Darwin’s Dangerous Idea* has an extensive discussion of design space.

[10] GRISOGONO, AM *A Generic Framework for Generating and Evaluating C2 Concepts*, 9th ICCRTS, Copenhagen, 2004 for discussion of the seven problems of transformation.

[11] RIDLEY, MARK *Mendel’s Demon: Gene Justice and the Complexity of Life* ?

[12] WAGNER, GÜNTER P AND ALTENBERG, LEE *Complex Adaptations and the evolution of Evolvability*, preprint

[13] SCHUSTER, PETER *Evolution in silico and in vitro: The RNA model* preprint