9th International Command and Control Research and Technology Symposium

Coalition Transformation: An Evolution of People, Processes and Technology to Enhance Interoperability

Track: C4ISR/C2 Architecture

WHAT DO NATURAL COMPLEX ADAPTIVE SYSTEMS TEACH US ABOUT CREATING A ROBUSTLY ADAPTIVE FORCE?

Anne-Marie Grisogono

Defence Science and Technology Organisation
PO Box 1500 Edinburgh 5111
South Australia

Ph +618 8259 6532
Fax +618 8259 5055
Anne-Marie.Grisogono@defence.gov.au

WHAT DO NATURAL COMPLEX ADAPTIVE SYSTEMS TEACH US ABOUT CREATING A ROBUSTLY ADAPTIVE FORCE?

Abstract

Wherever we look we find naturally arising Complex Adaptive Systems (CAS) in a wealth of forms and with an impressive display of variety and emergent properties, including the ability to adapt and thrive under pressure, change and competition. Of course there are failures too: species become extinct, individuals are defeated and die, stock markets crash, and brains develop mental illnesses. A defence force is also a CAS, and can exhibit both spectacular successes and failures under pressure, change and competition – but there is an important difference from naturally arising CAS such as organisms or ecosystems. The design and structure of a defence force is not evolved entirely through blind selection processes, but is significantly shaped by conscious human decisions which aspire to increase the 'fitness' of the force for some externally defined purpose. However, the very complexity of the decisions means that we rarely get the outcomes we want or expect. Yet on the other hand, we observe some natural CAS exhibiting both greater complexity and more of the robust adaptivity that we are seeking to build into our defence systems. This suggests that the natural world has much to teach us about dealing effectively with complexity.

Analysis of a range of adaptive mechanisms found in natural systems is a good place to start, and this leads us to a generic model of adaptation which can then be applied to engineered systems in a number of ways: firstly as a template for identifying the informal adaptive loops that arise spontaneously in any complex sociotechnical system, coexisting with, and often undermining the deliberate formal adaptive mechanisms; secondly for analysing the factors determining the effectiveness of both formal and informal adaptive mechanisms, and suggesting leverage points for modifying them; and thirdly, for designing new adaptive mechanisms to deal with anticipated future pressures and changes. All three of these approaches lead to significant insights into how defence systems should be architected, and how decisionmaking, C2 processes, information and policy should be organised and managed for the required degrees of robustness and adaptivity.

1. Introduction

Complex Adaptive Systems[1] (CAS) are a subset of Complex Systems[2] which exhibit the property of adaptation. They occur naturally throughout our environment, at multiple scales in living systems and in their products such as economic and cultural systems. A defining characteristic of a CAS is the ability, within limits, to change its composition and/or behaviour in ways that improve its overall ability to exploit its environment and thrive. Clearly this is something we would aspire to achieve for our defence forces, so it behoves us to understand the factors for success and failure in natural CAS and to explore resulting insights into how we may apply those factors to better purpose in the design and management of our defence capabilities.

We commence with an overview of the generic features of CAS and then turn our attention to a systematic analysis of adaptation in natural systems, identifying its principal features and mechanisms. This allows us to formulate a generic model of adaptation which can be applied to defence systems in a number of ways:

- firstly as a template for identifying the informal adaptive loops that arise spontaneously in any complex sociotechnical system, coexisting with, and often undermining the deliberate formal adaptive mechanisms;
- secondly for analysing the factors determining the effectiveness of both formal and informal adaptive mechanisms, and suggesting leverage points for modifying them; and

[1] The CAS terminology is most prominently promoted by John Holland (1995) and often used to describe the research agenda at the Santa Fe Institute. See the popular science book by Waldrop (1992) for an overview of this research.
[2] The field of complex systems is also known as nonlinear dynamical systems, and includes systems whose elements may not adapt such as the weather system or a chemical reaction. See Bar-Yam, Yaneer. 1997. *Dynamics of Complex Systems*, Perseus, Reading, MA. for an introduction to complex systems.

- thirdly, for designing new adaptive mechanisms to deal with anticipated future pressures and changes.

All three of these approaches lead to significant insights into how defence systems should be architected, and how decisionmaking, C2 processes, information and policy should be organised and managed for the required degrees of robustness and adaptivity.

We will conclude with a brief overview of the current research agenda and an update on progress to date.

2. Features of Complex Systems

Complex systems consist of large numbers of simple interacting agents, and tend to have hierarchical structures where the emergent entities at each successive level are formed from coherent groupings of interacting elements at the level below. As an example subatomic particles can form interacting patterns which we recognise as atoms, interacting atoms form molecules, interacting molecules exhibit a wide range of chemical properties, and so on.

Emergence is often described as the appearance of surprising or unpredictable properties of a complex system, but a more accurate definition is the appearance of a phenomenon at one level of structure, requiring a new concept to describe it which is simply not applicable or meaningful at the level below, for example valence has no meaning for a subatomic particle.

Emergence is in fact a key property that distinguishes complex systems from merely complicated systems. While both kinds of systems may consist of numerous interacting components, a complicated system is characterised by a low behavioural complexity: systems such as a mechanical clock only perform a limited function, and if the interactions between components are altered the system "breaks", and typically no longer performs a useful function[3]. In contrast, complex systems have a high behavioural complexity and may be difficult or impossible to predict under some conditions, while displaying a high degree of stability in so-called 'attractor' states under other conditions.

The interactions in complex systems have some generic properties. Viewed as a network with interactions as edges and the system components as nodes, real complex system turn out to be neither aggregates of independent nodes nor fully connected. Kauffman[4] has shown that systems that exhibit complexity have an intermediate degree of connectivity between ordered (sparsely connected) and chaotic (highly connected) systems. From the characteristic that networks are not fully connected we can deduce another property of complex systems: interactions tend to be local and short range, so information propagates by rumour-like mechanisms rather than global broadcast, and the result of interactions are non-linear, which means that linear superposition cannot be used to derive system properties.

3. Features of Complex Adaptive Systems

In addition to having all the properties of complex systems, CAS have two key properties that allow the emergence of adaptation: diversity and internal models. Diversity means that a CAS has a set of alternative system behaviours, and implicitly requires a selection mechanism that can select from among this set. Without diversity, there would be no variation in system behaviour and no available options for changing behaviour over time. Diversity is the necessary raw material for adaptation.

Internal models are a representation within the system of the external environment. Internal models do not need to be explicit, but as a minimum should satisfy two requirements. Firstly, there must be systematic covariance between a property of the environment and a property of the internal model. For example, a sundial represents change in time by a movement in the shadow of its pointer across a numbered plate. Secondly, the internal model must influence the system's

[3] This implies we can describe the behaviour of the "broken" system states with a small amount of information, thus its behavioural complexity is low.

[4] Kauffman, Stuart. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford Univ. Press

behaviour in some way. The sundial by itself is not an internal model. However, in a system including the sundial and a person who uses the sundial to decide when to perform an activity, the sundial acts as an internal model (internal to the system, not the person). An internal model allows the behaviour of a CAS to be sensitive to its environment.

Complex adaptive systems come in many guises, and tend to spawn other complex adaptive systems. Living organisms are a good example. They not only consist of many levels of complexity, the levels ranging from biochemicals, through cycles of biochemical interactions, subcellular machinery such as ribosomes and mitochondria, cells, tissues, organs, functional systems such as circulation, digestion etc, the whole organism, local populations, species, and ecosystems, but they also exhibit adaptation at every level, and through the emergence of culture in populations of human beings, have spawned similar hierarchies of CAS in political systems, economic systems, and so on.

In addition to diversity and internal models, Holland[5] lists aggregation, building blocks, nonlinearity, tagging, and flows as characteristics of CAS[6]. Nonlinearity and flows (interactions) are more general properties of complex systems discussed above, while aggregation, building blocks and tagging are common, but not essential features of CAS.

Successful CAS display an exquisite and complex balance in that they are both sensitive and insensitive to small disturbances, in different ways. The insensitivity is what gives the attractors their stability, and the emergent properties their persistence, what we might call robustness in military terminology, meaning that quite a few things can go wrong in quite a few ways before functionality is seriously affected. But it is also true that small disturbances (or noise) are vital for adaptation to occur at all, since they seed variation. In that sense a small disturbance can lead to a large scale change in the CAS rather than quickly petering out.

The trick that CAS pull off is the balance between just enough robustness and just enough changeability. This results from the interplay of two opposing types of mechanisms, positive feedback loops that drive growth and change and amplify the small disturbances, and negative feedback loops that exert restorative control to the attractor states, creating stability. The interplay is critical, since too much of one or the other results in either 'eternal boiling' or total stagnation. The way they are linked is critical too, since disturbances need to invoke both kinds of mechanisms in just the right kind of way. Such intricate balance can only evolve over many generations of selection and we postulate that much of the evolutionary adaptation of CAS has to do with tuning these feedback loops and the links between them.

CAS are not controlled by any single element of the system. The coherent behaviour which they display arises through distributed processes of competition and cooperation between elements rather than by centralised control.

Finally CAS are open systems – they exist in a context with which they continuously interact through porous boundaries. There is generally exchange of energy, information and physical resources between the CAS and its context. The laws of thermodynamics tell us that the order exhibited by CAS comes at a price: they must export their waste heat and entropy to their environment and import higher grade energy and information, else their order and structure will soon decay.

4. Adaptation in naturally arising Complex Adaptive Systems

Adaptation is manifested in widely varying forms in different CAS, but there are underlying characteristics common to all of them. We are interested here in extracting the essence of adaptation and understanding the factors that determine the extent and robustness of its success. So as a first step we outline a generic model of adaptation and then go on to discuss and

[5] Holland, John. 1995 *Hidden Order: How adaptation builds complexity*. Cambridge: Perseus Books.
[6] These are discussed in detail in Grisogono, A., Ryan, A.J. *Designing Complex Adaptive Systems for Defence*, Systems Engineering Test and Evaluation Conference, 2003 Canberra.

compare its instantiation in natural CAS at two different levels of complexity: an organism and a species.

The first observation to make about adaptation is that it presupposes a concept of ‘better’ or ‘fitness’, which must have something to do with the context outside of the system itself. We will take this as a given for the moment and return to discuss it later.

The second observation is that the ability to change is a precondition for adaptivity, but although complex systems can exhibit dynamic behaviour patterns which may appear to be changing in an ordered way, such changes are not necessarily adaptive – they may be simply generated by the built-in rules of interaction between the elements. What is interesting is when the changes actually are adaptive, and the system starts to encode information about the context.

Figure 1 shows the essentials. Assuming for the moment that the context is fixed, suppose that as a result of an inner change in the system (the ‘variation’), its next interaction with the context produces a different outcome in the system (the ‘feedback’), and that different outcome happens to have the effect of either reinforcing or inhibiting what changed in the system (the ‘selection’). This is not quite adaptation yet, because ‘good’ changes might get inhibited or ‘harmful’ ones reinforced, but at least the system is responding to the environment in a way. The final step needed is for the selection process to be linked to fitness so that ‘good’ changes are preferentially reinforced and ‘bad’ ones inhibited. In that case the system is adaptive - its fitness will improve over time. We would say it is ‘learning’.

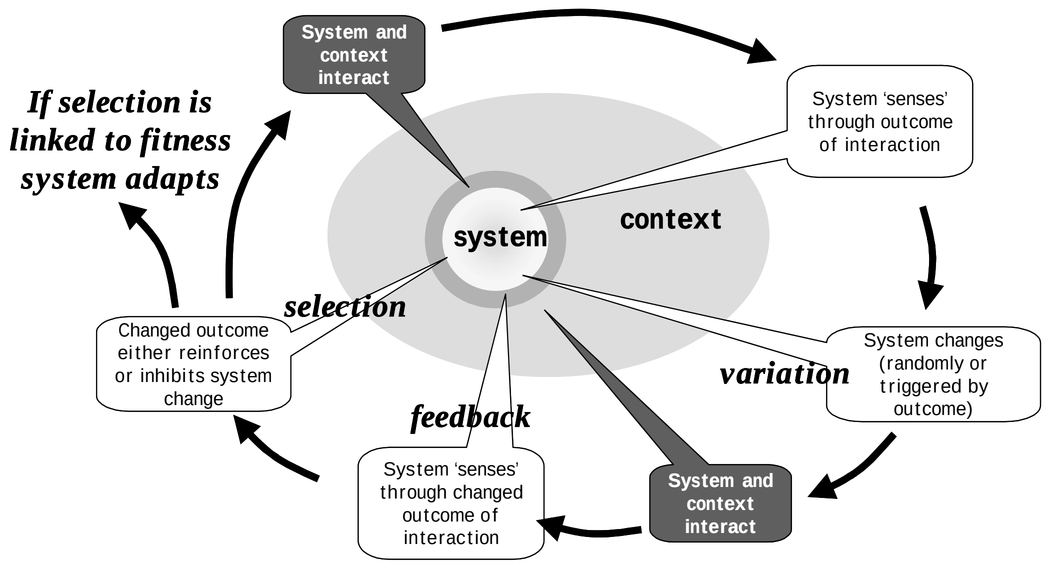

Figure 1 *Key features of adaptation: Variation of system, feedback through interaction with context, and selection of adaptive system change by fitness-linked reinforcement or inhibition. The selection process may become adaptive through evolution. This is a fixed context, with system fitness improving over time*

Such a mechanism acts as a ratchet, selecting for retention those variations that increase fitness and discarding those that decrease it, so that overall fitness is gradually increased. So through interaction and the adaptive process, information about the context is extracted from the environment and encoded into the pattern of fitness-enhancing variations that are retained, and as a result, the system becomes better able to thrive in that environment.

Interestingly, adaptivity at the level of the individual system can evolve even if it wasn’t there to begin with. Of course evolution is itself an instance of adaptation, but now at a higher level of complexity, the level of a population of such systems. If there is variation in the population with respect to how the selection mechanisms are correlated with fitness (say from total covariance through to total contravariance) then those that happen to have the more covariant selection mechanisms will thereby develop higher fitness while those that landed in the contravariant end of the distribution will become less fit. In the presence of competitive replication (i.e. systems that are 'fitter' are more likely to replicate) then over a few generations the proportion of systems having higher levels of fitness increases, and what was initially a random variation gradually becomes a feature of that population. By the same process, the other details (variation, interaction and feedback) of the adaptive mechanism can also evolve over time to become more and more effective in improving fitness.

If we now allow the context to change as well, as in Figure 2, the process works in exactly the same way, the only difference being that the changed outcome 'sensed' by the system is partly due to the change in the context. This may result in an increase or decrease in fitness, but provided there is still a source of variation in the system's interaction properties and the selection mechanism has sufficient covariance with fitness the system's fitness will improve nevertheless. As a result, the system is able to track changes in the context and maintain fitness, but only if the time constant of change in the context is much slower than the time constant for the system's response. If they are comparable the system may go into oscillatory behaviour, while too rapid changes in the context will induce chaotic behaviour and rapid collapse of the systems.

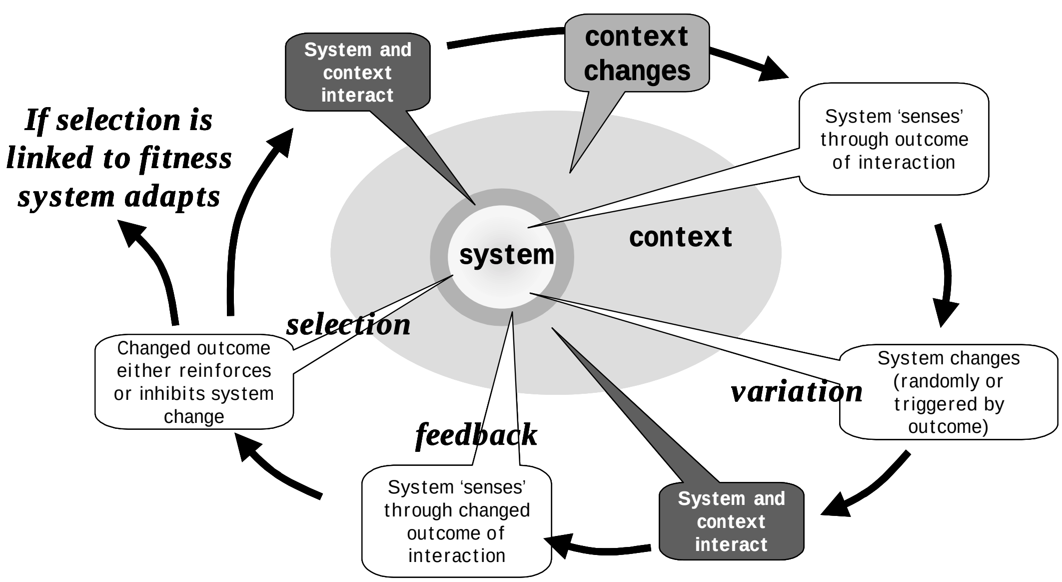

Figure 2 *Adaptation enables systems to respond to change in the environment. The system can track changes in its environment and remain adapted provided* $\tau_{context} >> \tau_{adaptation}$

We observe that there are some key differences between how adaptation works at the level of an individual system or organism (learning) and at the level of a population (evolution).

The most obvious one is the degree of parallelism in the process. A single system adapts through a sequential process where one thing is tried at a time, whereas a population is a parallel processor, each member of the population simultaneously trying one of a range of possibilities. But although evolution can 'evaluate' large numbers of possibilities at once it takes a whole generation for just one round of modifications, whereas a single system can try a number o fthings in rapid succession.

An important difference between individual adaptation and evolution in a population is in what constitutes fitness, and what the outcome of the process is.

For an individual system, fitness is the probability that it thrives. The outcome is that it becomes more responsive to its own local environment through encoding relevant information in its decision mechanisms, but the actual design of the system is not affected. We could say that learning improves how a fixed design is used in a given context. However the gain that learning produces is local and confined to the individual[7].

For an evolving population, the fitness is the probability that the individuals survive so that their genes get passed on into the next generation, and the outcome is that the population becomes more responsive to large scale changes in the context through encoding information about the context into the actual design of its constituent systems. We could say that evolution improves the design in a changing context. The gains resulting from evolution spread throughout the population and propagate in time through generations as long as they remain adaptive.

However, we note the limitation imposed on evolution's ability to generate complex new design features over a number of generations by the requirement that not only must each intervening generation be viable, but the successive forms of the new feature must be adaptive at each step. What this means is that evolution can only explore the neighbourhood (in an imaginary 'design space') of the current system design and that each point along the path has to be successful enough for the bearers of that genoytype to survive and reproduce so that it appears in the next generation.

Another key difference is in what gets retained or discarded.

In the case of an evolving population it is the combinations of characteristics embodied in its members – this particular combination dies out without being replicated and disappears from the population, that one is passed on through many descendants and becomes more frequent in the next generation. So there are two quite distincts aspects that evolution modifies, one is the spectrum of combinations of characteristics (genotypes) that exist within the population, the other is the distribution of frequencies of those genotypes. Together, these constitute an encoding of information about the context, and that information resides in the population of genotypes.

In the case of an individual system it is what we might call strategies that the system employs in its interaction with its context that get reinforced or inhibited. But there is more to it than just a collection of strategies – the overall behaviour of the system must include whatever process determines which of its strategies gets invoked, and under what circumstances. We can think of the strategies as a collection of options together with propensities for using them, and those propensities might be conditional on a set of system-dependent and context-dependent parameters, i.e. information about both the state of the system itself and about the environment, which together we can think of as a set of stimuli. In other words by paying attention to more relevant stimuli, having more effective strategies at its disposal, and more correctly associating the stimuli with the appropriate choice of strategy, a higher fitness system will be more likely to behave successfully in the particular circumstance.

So in this case there are three aspects for adaptation to work on: which strategies exist in the collection, what constitutes the stimuli – i.e. what information from the system and from the environment is 'used', and how the system associates the stimuli with the strategy options in order to invoke the one to be applied. Now the information about the context is contained in all three subclasses of information, and is physically encoded in whatever subsystems reside between the sensing and the acting system functions – the central nervous system in the case of a living organism.

We note the qualitative difference between learning, where there are three subclasses of information, and evolution, where there are only two. The strategies and their propensities (minus

[7] Transmission of learning between individuals is an example of a higher order adaptation that occurs in a society of individuals when culture emerges – we will discuss this level of adaptation, and its consequences for defence in a future paper.

the context sensitivity) in the learning case correspond to the genotypes and their frequencies in the evolution case. The odd one out is the stimuli – information used to assess current context in order to make the choice of strategy context sensitive – this has no counterpart in the case of evolution because there is no mechanism for evolution to ‘decide’ which genotypes it is going to try – it just tries them all at once blindly. So the context sensitivity of learned strategies can be seen as a compensating mechanism for the lack of highly parallel processing compared to evolutionary adaptation.

The encoding of information about the context raises an interesting question as to where the information comes from. We noted above that adaptation acts as a ratchet to extract information from a stream of random inputs, the variation in the system. That sounds magical until we realise that the information enters the system through the selection process (this one works, that one doesn’t…), not through the random material that selection works on.

Nevertheless, the variation is where the diversity on which adaptation thrives, arises and it is profitable to examine where the raw material comes from. In the case of evolution, there are four categories of mechanisms that result in changing the content of the genotype: random mutations induced by radiation or chemical reactions; recombination in the process of sexual reproduction; copying errors; and retroviruses.

Interestingly enough, there is a close parallel with the possible sources of raw material for trial in the case of learning. Random mutations amount to trial and error, a well-known technique for finding new and betetr ways of doing things. Recombination corresponds to the creative process of trying new combinations of elements of previously successful strategies that worked in different situations. Copying errors arise when previously successful strategies are misremembered. But what would retroviruses correspond to? A form of learning that hasn’t arisen yet, where a parasitic lifeform directly injects its own strategies into the organism’s central nervous systems.

Another important set of differences arise when we examine how interaction and feedback operate in each case.

For a species, the context is the physical environment and the other species it interacts with in its ecosystem, and it ‘senses’ its context through the many interaction of its various subpopulations of genotypes utilising their functionalities to perform the tasks required for survival and reproduction. Some genotypes are more successful than others, one could say the species experiences different consequent success rates in those tasks for different genotypes. These different survival and reproductive success rates directly modify the frequency distributions of genotypes in the next generations – a very direct form of feedback, so the selection process (what gets retained, what gets discarded) is tautological since reproductive success is both the measure of fitness and the selection maechanism.

For an individual, its context is the physical environment, other members of its own species, and members of other species. In the course of performing the functions necessary for survival and reproduction, it interacts with its context in a variety of ways, for example through predation, competition or cooperation. As each situation presents itself, the individual senses some aspects of the situation (the stimuli it has sensitivity to) and through some kind of processing (which we will describe without explanation as a neural activation pattern) it chooses a response strategy (eat, run away, snarl…) from its range of options. Its action has consequences of course (satisfaction, avoidance of a costly fight, success in warding off a rival…), and as a result of sensing those consequences (as another neural activation pattern) the original neural activation pattern that led to the choice of that strategy in the presence of those stimuli gets either reinforced or inhibited. This latter is the selection part of the adaptive mechanism, but note the qualitative difference from the evolutionary case: the consequences of action don’t directly impact on the retention or otherwise of those strategies. It is only through an intermediate mechanism of a consequent neural activation pattern having an inhibitory or reinforcing effect on the neural activation pattern that led to choosing the strategy that the organism can adapt its behaviour. The existence of an intermediate mechanism weakens the linkage between trial and learning, but is necessary in the absence of a more direct linkage, since fitness is here quite distinct from the selection mechanism. In higher organisms the reinforcement or inhibition is experiences as pain or gratification, what we might colloquially describe as the sticks and carrots approach. In any case, the consequence of successful learning is that the organism is more likely to choose an appropriate strategy in the future.

The choosing function is also radically different in evolution and in learning. An organism, as we have just discussed, will do a certain amount of internal processing of its informational inputs resulting in a choice, although that processing could be as simple as a biochemically encoded stimulus-response mechanism, or as complex as a human imagining the consequences of various options. The latter is an example of pre-selection of likely successful strategies by running an internal model of how the interaction will proceed. In the case of evolution however, the notion of choosing is almost irrelevant – which genotypes arise in each generation is largely a matter of chance – but there is one factor operating which resembles choice: a crude filtering out of totally unworkable genotypes as non-viable embryos that fail to develop.

5. Adaptation in Engineered CAS

Understanding the mechanisms for adaptation in naturally arising complex adaptive systems suggests some insights into how designed systems could adapt at different levels.

The most important difference to observe when comparing natural CAS to designed CAS is that natural systems do not have any externally defined purpose or meaning, whereas defence systems and operations are designed for a specific purpose and the fitness measure for this purpose is artificially imposed upon the system. It is therefore imperative that the selection measures and feedback mechanisms contain strong links to the fitness function if adaptation is to drive system evolution in accordance with the system's intended purpose.

We can apply the same framework of analysis to examine instances of, and opportunities for adaptation in engineered or designed systems.

We postulate two kinds of engineered adaptation corresponding to the two naturally arising ones discussed in the previous section; an adaptive system, and an adaptive force, and discuss each one in turn.

For an adaptive (learning) system, the fitness (selection mechanism) is essentially the probability that the designer thinks the system will achieve its intended purpose.

The encoded information resides in the system between its sensing and acting subsystems, in the form of data or system parameters.

The choice of response in a given situation may be based on either direct feedback from interaction with context, or from preselection of likely 'winners' and avoidance of 'killers' by internal modelling of the anticipated interaction with context, based on sensed information.

The selection mechanisms for retention or discard of strategies may include, depending on its design, the ability to use success or failure data from the context to modify any of: its sensed inputs, its response options or its response-choosing algorithms.

The action options possessed by the system may be fixed or open, and if they are open, the source of new options to try may come from random trial and error, from algorithms operating on constrained inputs, or lookup tables operating on a fixed menu, or through a creative type of process exploiting recombination, neural nets, or search algorithms.

The outcome of implementing learning-type adaptation into a system is that it becomes more responsive to its local environment, and improves how its fixed design is used, however while the gain is still local as it was for the natural counterpart, it is more easily shareable.

The parallels with the natural process of learning are clearly very close, but there are some interesting differences and opportunities.

The main insight we offer here is that the even more remote connection between the 'fitness' implied by how selection is implemented and what the fitness is actually intended to be, poses real difficulties for successful exploitation of learning in engineered systems.

Another insight springs from the observation of the three subclasses of information encoding in the natural learning case. We suspect these have not been widely recognised as such, nor exploited in the design of learning systems. The three suclasses correspond to allowing learning modifications in what information is sensed and utilised, in what response options are generated, and in how the two are linked or associated.

This is related to Ashby's law of Requisite Variety which tells us that there needs to be a commensurate degree of choice in the actions that a system can perform with the number of states that its control system can process, but here we see that there is an informational aspect as well which has to match. This perspective suggests a much richer range of options for developing learning systems than simply weakening or strengthening the choices form a fixed menu in the light of experience.

The second example we have chosen, corresponding to evolution, is the context of force development. In this case the information encoded resides in the architecture of the system-of-systems comprising the force, and in the technical information about its elements, in its processes, and in its organisational structure. The fitness, based on what selection mechanisms actually operate, is the probability that the decisionmakers think that the development will achieve the required intent, or perhaps more accurately, the probability that a particular development option is implemented.

Adaptation at this level ought to be parallel if it is to follow the form of natural evolution, however it is parallel only at low end of small innovations arising from the large number of defence personnel simultaneously trying to solve their own local problems. At the expensive end of capability debelopment, the process is excruciatingly serial – the processes of strategic analysis, systems engineering, system integration, acquisition, doctrine development, plus the impact of political, sociological and economic factors.

Interaction with the defence context and within system occurs via competition for limited resources required for implementation of defence development options. Active selection is based on modelling to preselect likely winners but many other factors, not necessarily adaptive, come into play.

The source of variation is scientific research; human creativity; computer-aided design, analysis of recorded information, and the overall outcome is that the force becomes more responsive to large scale changes in environment, and explores the design space neighbourhood in order to produce new design features.

The chief insight this comparison yields is once again the serious disconnect between the selection mechanism operating and the intended fitness.

This is the result of the added complication of a conscious human decisionmaker who may have an idea of what the 'fitness function' is intended to be, but a very incomplete and inadequate concept of the consequences of the decision options on what firness will actually result.

Another aspect that comes into focus when we examine defence systems from a perspective of having looked at natural complex adaptive systems, is that defence already exhibits many of the features of a spontaneous natural CAS. This should not be unexpected, consisting as it does of a large number of elements, interacting in different ways, and with hierarchies, tags, specialised roles, systems of reward and punishment, and so on. There are large numbers of formal and informal positive and negative feedback loops operating and it is inevitable that these interact and produce unexpected behaviours.

Defence tends to be a very tightly controlled environment by its very nature, but it is also true that it is very conservative, and resistant to change. These properties can be understood as arising from negative feedback loops, commonly described as ‘culture’. When change is proposed from above, the attempts to implement it often produce mixed or disappointing results, and that is quite likely a result of interference between the formally imposed control mechanisms and the informal ones.

As a complex enterprise, defence may be more amenable to steering through a better understanding of those informal CAS mechanisms and deconfliction with the formal processes.

Axelrod and Cohen[8] have discussed a number of ways in which complexity can be harnessed in the management of large and complex enterprises, and many of them are directly applicable to defence.

We will discuss these opportunities further in a subsequent paper.

6. Conclusion

In summary, there are a number of factors contributing to the challenges of designing and implementing adaptation in defence systems

- Lack of clarity about what intended fitness is
- Measuring fitness/success in real world, estimating it in models
- Disconnect between fitness and selection mechanisms operating
- Deciding what to sense in the context i.e. what stimuli to pay attention to (the inputs)
- Generating new more successful responses (the outputs)
- Encoding of information about the world into the system, evaluating candidate responses and understanding under what conditions to pick which response (the links)
- Understanding the reinforcement and suppression mechanisms already operating and being able to modify them to support the intended adaptation
- Deconflicting self-organised and imposed mechanisms
- Being able to execute the adaptive cycle fast enough to keep pace with changing context
- The need to parallelise the selection process and the difficulty of doing so
- The need to forecast rapid changes in the context and the difficulty of doing so
- The need to quickly implement the adaptive changes identified and the difficulty of doing so
- How to handle the linkages between levels of adaptation at different levels of complexity – decoupling them for manageability, managing the coupling for creating conditions for success
- The difficulties of operating in an adaptive context (both internally and externally to own force) and the resulting danger of instabilities, the difficulties of managing the required co-evolution of capability elements..

These issues form a stimulating agenda for cross-disciplinary research between the fields of CAS and defence science and we believe that many powerful new results await discovery and exploitation.

---

[8] *Harnessing Complexity: Orgaizational Implications of a Scientific Frontier* R Axelrod and MD Cohen, 2001, Simon and Schuster.