

Complex Systems vs Complicated Systems

Glynn Rogers

CSIRO Telecommunications and Industrial Physics

A quick look at some of the literature suggests that the questions of what is a 'complex system' and what are 'emergent properties' are very much open and the subject of continuing debate. However there does appear to be some consensus that the context of the debate involves different levels of description - there may be multiple levels as in biology but, for the sake of simplicity, let us consider just two, the micro and macro levels.

If complex systems are to be regarded as consisting of multiple interacting entities then these entities exist at the microlevel and their properties, and their interaction mechanisms, are described in the microlanguage. In the case of nonlinear dynamical systems, phase space, trajectories, and the zoo of attractors, are all denizens of the microworld and the mathematics of coupled first order nonlinear differential equations forms part of the microlanguage.

If the machinations in the microworld cause the system as a whole to exhibit some form of structure - temporal regularities, spatial patterns etc - then this structure constitutes the macrolevel and its description uses the macrolanguage. The central feature of the complex systems concept is that this structure is not observable at the microlevel and cannot be described in (reduced to) the microlanguage.

Are we done? - does this resolve the problem? No - far from it I believe. Let me go back to my car example. The various components of the car - engine subsystems, steering subsystems, braking subsystems etc are each microlevel entities with their own properties described in the microlanguage of mechanical engineering and which interact to produce macrolevel behaviour - the dynamics of handling, the characteristics of 'ride' and so on. But, in spite of my engineering bias, my intuitions are uneasy at the idea of calling a car a complex, as distinct from a complicated, system.

There is clearly more to this than simply invoking different levels of description and this troublesome notion of emergence seems to be an essential part of it. The macrolevel structure should 'emerge' from the microworld whereas the macrolevel behaviour of a car seems to be merely a consequence of the microlevel properties. There is something deliciously mysterious about the notion of emergent properties - something which hints at new 'principles of organisation' etc.

Now I think this is largely illusion. There will of course be new interaction mechanisms to be discovered which will reveal new connections between features of the world, but I don't think 'emergence' is one of them - it just isn't that sort of beast.

Considerable light was thrown on this for me by a nice little philosophical essay in "Complexity" by Lee McIntyre who invokes the fundamental philosophical distinction between matters of ontology - what there is - and matters of epistemology - what and how we can claim to know. In his view emergence is an epistemic not an ontological notion - that is, it is a characteristic of the structure of our knowledge of the world, not something which is out there in it. I think he is right about this.

It seems to me that there are two consequences which are quite important for our task. The first is that 'emergence' and the distinction between 'complex' and 'complicated' lie, like beauty, in the eye of the beholder. That is to say we will not be able to distinguish between complex and complicated systems by

some form of empirical investigation to discover whether a given system does or does not have some property of 'emergence'. Rather it is a distinction which we can make by suitably structuring our knowledge of any particular system to give it the characteristics of a complex as distinct from a complicated system. Furthermore we need to decide, on utilitarian grounds, just what those characteristics are, for example that suitable redefinition of control parameters and variables can transform the system's description into a form which enables it to be designed to be largely self regulating or self assembling.

The second consequence, which actually follows from the first, is for the way in which we structure the complex systems activity in CSIRO. This should be done so as to provide a strong motivation for work in the various application areas to identify complicated systems as candidate complex systems and to structure their descriptions (and perhaps the systems themselves) so that they can be brought under the complex systems umbrella. Of course it is not currently clear, and probably never will be, just what that umbrella is, so this implies a considerable amount of interaction between the 'theoretical core' that we envisage and the application areas. Indeed I would expect this interaction to drive the evolution of our understanding of the nature of complex systems. A complicated system is then one which resists this effort.

The central point here is that we should not be attempting to set up some rigid set of admission criteria by which application areas can be categorised as either complex or complicated with only those passing the test being legitimate targets for complex systems activity. This is what I have tried to capture in the accompanying diagram. At the centre are the quite theoretical entities which provide the conceptual and mathematical tools to tackle Complex Systems Science. Surrounding this is a belt of idealised applications of the theoretical entities which illuminate them and which provides a 'matching layer' for the outer region of more messy (complicated) real world problems.

My basic point is that there is no well-defined boundary between the complex and the complicated - as we move outwards, the flavour of complexity decreases and the flavour of complicatedness increases. The pragmatic goal of Complex Systems Science, it seems to me, is to increase the flavour of complexity in the outer regions by a combination of theoretical analysis, empirical investigation and, importantly, problem restructuring and design. Achieving this will require the combined efforts of the domain experts out in the Divisions and the complex systems theoreticians in the Centre. Ensuring a vibrant interaction between the two groups (and between experts in different domains) in an organisation of 'medieval fiefdoms' is our challenge.

Let me illustrate these ideas with a couple of examples drawn from the world of nonlinear dynamical systems. Some of you may regard this as instruction in the fine art of egg sucking but bear with me as I run through the basics of the story to get the context right. We deal with a multivariate system described by a set of coupled, nonlinear differential equations so that the instantaneous state of the system is a point in a multidimensional phase space and the dynamics occurs along trajectories in that space. Assume that the system is at a stable equilibrium point, a fixed point attractor, and that the response of the system to small perturbations is analysed by linearizing the equations about that point. The real parts of the eigenvalues of the resulting matrix are negative so that the perturbations decay exponentially to zero with time.

Now assume that some small subset of these eigenvalues (say one) are much larger than the rest ie closer to zero, which is a very common situation. Then the perturbation will decay much more slowly along the corresponding eigenvectors ie the dynamics will be predominantly in a low dimensional subspace, perhaps along a single direction. The resulting dynamics, while involving all of the system variables, perhaps roughly equally, are low dimensional ie they form a simple structure or pattern. For example if

there are two dominant eigenvalues which are complex conjugates then the dynamics will be a decaying sinusoidal oscillation with frequency equal to the value of the imaginary parts.

The essential point here is that this dominance can be formally exploited, in what Herman Haken calls the Slaving Principle, to analyse the effects of changes in the control parameters on the stable point, particularly where some form of phase transition is approached. By transforming the system into a set of coordinates determined by the eigenvectors of the linearised equations, the new variables along the eigenvectors with small eigenvalues can be expressed in terms of (or slaved to) the variables associated with the dominant eigenvalues. The nonlinear analysis can then be carried through, at least for small changes, with a much smaller subset of variables - those corresponding to the dominant directions in the original linear analysis. Because the analysis requires only a small subset of variables, the dynamics associated with an attractor resulting from a transition from a precursor cannot have an arbitrary increase in complexity. There just isn't a way of getting enough information in there - simple patterns of behaviour beget simple patterns of behaviour.

An example is the Hopf bifurcation where the real parts of two complex conjugate eigenvalues become positive with changes to the control parameters and the original equilibrium point becomes unstable leading to the formation of a limit cycle attractor. Because the transition takes place predominantly along directions associated with the original conjugate pair, the resulting attractor retains the character of the original decaying oscillation and has a simple low dimensional behaviour pattern - a limit cycle.

So the Slaving Principle provides a relationship between the patterns observed in the macroworld and the detailed machinations in the microworld. It explains, in a sense, how the former can arise from the latter and, importantly, why we can get quite simple patterns or a high degree of 'self organisation'. Here, I think, is the basis of a distinction between complex and complicated - a complex system has a unified theory that explains how the structures which constitute the macrolevel arise from the microlevel. Complicated systems like a car do not have such a unified theory although they may have a set of theories linking specific parts of the microlevel to particular patterns of behaviour.

Is this Slaving Principle then a new principle of organisation that reveals 'emergence'? Well, not in an empirical sense - it was not an empirical discovery but the result of mathematical analysis. It is essentially a mathematical tool and there are presumably accounts of nonlinear systems analysis that don't use it, at least explicitly.

For example, in a recent paper in "Science", David Wales uses catastrophe theory to analyse changes in the Potential Energy Surface of atomic and molecular clusters as the nature of the cluster constituents changes. The PES describes the combined effects of the attractive and repulsive forces between constituents and is obtained by adding an extra dimension to the state space of the cluster. Catastrophe theory provides a means of analysing the behaviour of critical points in the surface as control parameters change in terms of a decomposition into elementary one dimensional polynomials. In particular Wales examines the coalescence of minima and transition points which is essentially what I was talking about above.

Now the nonlinear terms in a dynamical systems approach are just the gradients of this PES so that the matrix resulting from the standard linearisation procedure is just the Hessian of the potential whose eigendecomposition is used in the catastrophe theory analysis. So, presumably, (I haven't looked at it) the catastrophe theory approach is essentially equivalent to the Slaving Principle approach - it just looks at the issue from a different viewpoint.

Of course all this may seem to apply only to an small, rather theoretical, part of the world of interest only to physicists and our colleagues in molecular science. However a link to a much wider world is pointed to by a discussion of statistical self assembly by Tad Hogg from Xerox Parc in "Nanotechnology". He uses the notion of an 'energy' function to measure the degree to which a set of perhaps mutually inconsistent constraints are satisfied by a system of components self assembling via 'thermal' motion, into a structure represented by a global minimum of the function. (Presumably the mathematical basis of this is the Lyapunov function concept.) The salient point here is that this 'energy', as he points out, doesn't need a physical basis but could, for example, measure the degree of compatibility between computational elements and this, in turn, implies that the 'self assembly' process itself could be computational. He uses genetic algorithms and simulated annealing as examples.

Now this really is the thin end of the wedge because once we start going down this path there are immediate links to the world of cellular automata on the one hand and the galaxy of Bayesian analysis, via stochastic relaxation, on the other, particularly the Markov Random Field concepts used in image processing. I suspect there is much productive territory to be explored here. But there is perhaps a more dangerous path which leads to the whole world of nonlinear optimisation.

Let me illustrate the point by returning to my car example but this time considering the design and manufacturing process rather than the car itself. Currently cars are automatically manufactured, as distinct from self assembled - what Hogg calls directed manufacture - and the design process, presumably, still requires considerable engineering input. However suppose it was possible to structure the design and manufacturing process so that it could proceed from a detailed performance specification and a set of general constraints including those on the assembly process. The key would be a global nonlinear 'energy' function which measured the degree to which the specifications and constraints were satisfied. Suppose the end result of some optimisation process was then used to program the automatic factory so that the car which finally popped off the end of the production line was the result of a completely automatic process of detailed design and manufacture. (Various computer aided design systems used in electronics for example and the Boeing 777 indicate this is not an unrealistic scenario.)

What is the essential difference between this and self assembly (or the chemical processes considered by Wales for that matter)? My central point is that there isn't one - at least not in the sense that it exists as a brute fact about the world which we need to discover. Its up to us to draw any distinctions necessary on utilitarian grounds.

What are these grounds? It seems to me that we need to keep our own macro level clearly in focus here and ask what we are trying to achieve with Complex Systems Science. My answer, with my engineering bias, is that we are aiming to develop the techniques and process for designing and managing large scale systems of nonlinearly interacting components. In this sense, I see CSS as analogous to Engineering Science in that it draws upon the underlying 'pure' sciences to provide the tools to do this and, like Engineering Science, there are no firm boundaries between it and those sciences. And like Engineering Science, CSS should include reformulation of the systems it deals with in order to make them amenable to these techniques and processes.