

Understanding Observed Complex Systems

– *the hard complexity problem*

Bruce Edmonds

*Centre for Policy Modelling, Manchester Metropolitan University
Aytoun Building, Aytoun Street,
Manchester M1 3GH, UK.*

+44 161 247 6479

bruce@edmonds.name

<http://bruce.edmonds.name>

Abstract. Two kinds of problem are distinguished: the first of finding processes which produce complex outcomes from the interaction of simple parts, and the second of finding which process resulted in an observed complex outcome. The former I call the *easy complexity problem* and the later the *hard complexity problem*. It is often assumed that progress with the easy problem will aid process with the hard problem. However this assumes that the “reverse engineering” problem, of determining the process from the outcomes is feasible. Taking a couple of simple models of reverse engineering, I show that this task is infeasible in the general case. Hence it cannot be assumed that reverse engineering is possible, and hence that *most of the time* progress on the easy problem will not help with the hard problem unless there are special properties of a particular set of processes that make it feasible. Assuming that complexity science is not merely an academic “game” and given the analysis of this paper, some criteria for the kinds of paper that have a reasonable chance of being eventually useful for understanding observed complex systems are outlined. Many complexity papers do not fare well against these criteria.

Keywords: complexity, emergence, computability, reverse engineering, application, observed systems, methodology, critique

Introduction

It has now been repeatedly shown that producing complex phenomena from the interaction of fairly simple parts is possible. This is what I call the *easy complexity problem* because, now people have focused on emergent phenomena, it turns out not to be *too* difficult to achieve. These phenomena may be very interesting in the same sense that a puzzle or computer game is interesting, but the point of complexity science is that it is more than just interesting in this sense but has some exterior value.

What is harder is producing *significant* emergent behaviour, but this raises the question of what makes some emergent behaviour significant. I suggest there are

two criteria for such significance: *firstly*, that the emergent behaviour corresponds to some that is observed in natural systems or, *secondly*, that this behaviour has some generality (that is, corresponds to a wide range of possible behaviours). In a sense the latter ultimately reduces to the former because its generality is in anticipation that it will, in the future, be useful in understanding observed systems¹. If it turned out that it never, for some reason, corresponded to any observed system its interest would be, literally, academic.

This points to the more important problem facing complexity science, namely that of understanding an observed complex system. This is much harder because one does not merely have to find a set-up that, when interacting, results in a set of emergent phenomena, but needs to be the set-up that results in *this particular* complex behaviour. In other words, the hard problem that faces complexity science is how to “add value” to the understanding of *observed* complex systems. Making progress with the hard problem is ultimately much more important than academic progress with the easy problem.

Here it is argued that progress with the easy problem can only contribute a little (and only then if they are *very* lucky) to the hard problem, due to the remaining barriers between such ‘easy’ results and understanding observed complex systems. This is a result of the multiple difficulties in matching emergent outcomes to what is observed. These difficulties are for several reasons and are the core substance of this paper. Much complexity science tackles the easy complexity problem, only *implying* that it might be helpful in tackling the hard problem. However I argue that to be *actually* helpful with the hard problem² a change of approach is required.

Finally some criteria that might be used to judge whether some work in complexity science is actually helpful (or likely to be helpful) with respect to the hard problem are put forward. By these criteria many existing papers in complexity science do not fare well.

¹ Or in understanding further general patterns which are useful in understanding observed systems etc. The point being that if the chain of understanding *never* has any prospect in understanding the observed then it is ultimately sterile.

² To be precise: have any significant chance that it will eventually contribute to the hard problem.

The Difficulty of Reverse Engineering Complex Systems

The “reverse engineering” problem is: *given* a particular set of observed outcomes, *which* candidate setup could have produced this. Thus if you have a particular set of observed statistics about a certain outcome that is your target, one might be looking for the particular simulation that would match these (to a predetermined acceptable level) when run. Here the outcomes are known but the generating mechanism is sought for. This is the opposite of most papers in complexity science where the generating mechanism is the starting point and it is the outcomes that are investigated. In order for progress with the *easy problem* to be helpful with the *hard problem* then one needs to be able to reverse engineer the observed complex phenomena to work out which model of complex phenomena explains it – only then do the *easy* results become useful in sorting out the *hard* problem. Thus the core of this paper considers this reverse engineering.

If one had a theory that would help one match which set-up would result in which set of outcomes which was *reverse soluble*, in the sense that the right set-up can be deduced from the outcomes (for example by analytically solving some equations) then solving this problem is clearly possible. However for most of the processes that complexity science is interested in, there is no such possibility. For example in all “artificial life” processes that are *weakly emergent* (Bedau ????) then there is no way of relating the set-up and the outcomes other than by doing a simulation.

Clearly, the development of theories, however partial, that would guide us as to which setups might have generated any particular set of outcomes might allow some of the many results concerning the easy problem to be applied to observed outcomes. However, there is more than one difficulty in trying to do such reverse engineering. This section will look at some of these difficulties in the general case.

Consider a system that is weakly emergent, so that it is infeasible to reverse solve which setup caused the target outcomes. In this case there is some “search” process where one tries various setups in order to find those that match the outcomes to an acceptable degree. The difficulties of such a process could include the following sub-problems:

1. The search necessary for finding such a set-up is infeasible;
2. In the presence of noise in the core process, even the right set-up may not reliably result in the target outcome (i.e. only does so sometimes);
3. Due to noise in the output signal, the outcomes may only approximately match those of the target, leaving the problem of deciding when a match is acceptably close;
4. There may be more than one set-up that matches any given set of target outcomes to any given extent (including perfectly);
5. It is possible that no set-up might result in an acceptable match to them in all respects, but only with partial matches or no match at all.

The presence of noise does make this process more difficult, especially because slightly different noise can result in very different outcomes (for example, in chaotic systems). However some problems are not essentially due to the noise, it is that the noise is an *additional* source of difficulty. Thus first we consider the case in a model of the reverse engineering problem no noise, then move on to a model of the case with noise.

Without Noise

To show the difficulty of the reverse engineering problem *even in a noiseless environment* I consider a very general but abstract model of this problem. This is the task of finding a computer program, P , that reproduces a given infinite sequence of binary digits, T (for target). The program represents the complex process and the sequence of digits the outcomes of that process as it progresses over time. Clearly given the Church-Turing Thesis then any effective process (including any simulation that can be done on a computer) can be mapped onto this in a general and effective way.

Given this model, standard results of recursion function theory (e.g. Cutland 1980) tell us that:

- There are infinite sequences, T , which are not produced by any program P
- If there *is* a program, P , that produces T then there are an infinite number of other programs that also produce T
- There is no general and effective way (i.e. computable method) of finding a program P that produces a given T

- There is no general and effective way of checking whether a given program P produces a given T without running it forever
- There is no general and effective way of checking if two programs result in the same output T
- Even if T is produced by an (unknown) program there is *still* no general or effective way of checking that a given P does (without running it forever)

These results mean that sub-problems 1, 4 and 5 listed above are insoluble in general. Thus even in the absence of noise and even in the situation where the setups, the output and the process are utterly clear and known they are still insoluble in general. This does not bode well for a project that tries to apply results from the easy complexity problem to the hard one. We now look at a similar model of the case with some noise.

With Noise

There are several *ways* of adding noise into the above model (that of determining which program P can have produced a given T). There could be noise in the description of the program itself, similar to a mutation in an genetic algorithm; the process of computation could be prone to a level of error (like electrical noise in a processor chip); there could be explicit random elements in the program (like a random seed); and there could be a measurement error in comparing the output of P to T .

Here we look at the situation where there is the equivalent of a random seed, which is often interpreted as some randomness or uncertainty in the process itself. This corresponds to the case where P has a parameter which represents this noise and whose value is unknown to us who are considering its general properties, $P(n)$. This corresponds to a situation where there is a pseudo-random generator in the program. The n can be seen as an unknown input of the program or the random seed used in many simulations.

Now we have added some analogue of noise the situation with respect to the sub-problems is even worse, because not only do the difficulties with P (rather than $P(n)$) hold but there are new difficulties, including that:

- There is no general and effective method for determining the range of output sequences that $P(n)$ would produce for different n

- If there is an n so that $P(n)$ produces T , then there is no general and effective method for determining how different (measured say by hamming distance/length) the results of $P(m)$ will be, for a given different m
- If there is an n so that $P(n)$ produces T , then there is no general and effective method for checking if $P(m)$ produces T for all m

These difficulties mean that the sub-problems 2 and 3 are also insurmountable in general.

Feasibility and Scope

Thus in the general case, the outlook for the reverse engineering is bleak, and hence, in general, solutions to the *easy problem* cannot be assumed to be useful in solving the *hard problem* in any particular case.

This is a general and abstract result, that simply points out that there are some *very hard* cases out there. Of course, there may be additional knowledge about the kind of systems that are being dealt with that restrict the range of systems to those where the reverse engineering problem *is* feasible. However then, specifying what range of system is being considered (alternatively making explicit what *a priori* assumptions are being made about the systems being observed) becomes important, along with some justification why exhibited results that address the *easy problem* might be, in practice, useful with respect to the *hard problem* in the specific range of cases being discussed.

This is very evident in the extreme case, where the range of possible systems is restricted down to a very small range, as happens with detailed evidence-based simulation models. Here there are some unknowns, both assumptions about mechanism and parameter values, but these might be known to some degree due to fallible or partial evidence. If the outcomes are known sufficiently, especially if short-range or detailed outcomes are known as well as general properties of final outcomes, then it may well be possible to make some tentative conclusions concerning which set-up within the restricted range corresponds to the outcomes.

Thus what the above abstract argument concerning P and T does is to throw the burden of proof onto those that simply *assume* that solutions to cases of the *easy problem* will be useful with respect to the hard problem.

Discussion

Of course, one can always simply *assume* that any present paper might be useful in the distant future regardless of any present judgement of its ultimate relevancy and thus protect any particular paper against the charge of irrelevancy to actual problems laid out in this paper. This argument says that since, in the past, papers and ideas that were not thought to be useful turned out to be that the papers criticised here might also be. However this argument is spurious since it would support *any* paper against *almost any* criticism based on relevancy, however tenuous its chances of contributing, ultimately, to the *hard problem*. Clearly we have to make *some* judgements as to which papers should be judged as worthy of being read by others in the field and so necessarily there will be some criteria applied – the only argument is *what* those criteria might be.

In a sense what we need (without being too dogmatic about it) is to apply criteria that selects those papers that are *most likely* to be useful in the future. My argument (above) is that papers that tackle the *easy complexity problem* will be of limited use, unless there is some particular reason to think otherwise, given a particular restriction or property of the class of process being considered. In other words a *special pleading* why progress with the reverse engineering problem might be soluble in this set of cases, and hence that results with the easy problem

Some Criteria for Judging Complexity Papers

Clearly there are different ways in which a complexity paper, presented now could turn out to be useful, ultimately, with the *hard complexity problem*. Of course it could be argued (indeed has been argued, e.g. Feyerabend ????) that *any* prescription or normative judgement concerning the direction or approach of a field is counter-productive and that the *only* thing that works is “anything goes”. This would correspond to a “random search” in machine learning where no clues as to the direction of the solution are heeded but random solutions chosen and checked. However this is clearly not the case presently in complexity since papers *are* selected by peer review in conferences and journals as to their relevance and importance (as well as other important factors that are exterior to this discussion such as readability and soundness) – the only thing in question is *what* sort of relevance and importance criteria should be applied for the ultimate

and general success of the field. At the moment this seems to involve the extent to which a paper is interesting as a surprise or academic puzzle *under the unquestioned assumption* that such easy results will be ultimately useful for such as the *hard complexity problem*. I am arguing that this assumption is misguided and hence the balance of judgement should shift towards the relevance of results to the *hard* problem and away from those that only deal with the *easy* problem.

Hence I now suggest some criteria, each of which describes a different kinds of paper that has some chance of being ultimately useful with respect to observed complex systems – ones we actually come across in the business of life and technology. These criteria come out of the discussion in this paper. They include those that satisfy the following criteria.

- Pseudo Mathematics – It is clearly the case that abstract mathematics *has* sometimes turned out to be very useful despite its lack of direct relevance. It is also the case that some simulation work seeks to discover the properties of certain systems where analytic solutions are infeasible. However there are tough criteria for pure mathematics, those of importance, generality and soundness. If a paper claims aims to do something similar then these criteria should hold. If the results are important, general and sound then there is a chance that it may turn out to be useful for a future hard complexity problem, but if the results are limited to a particular set-up, small range of parameters etc. then this would fail under this criterion. In this case the burden of justification is to show that their explorations are sound, general and important.
- Following the argument in this paper *if* one is dealing with a restricted class of system where the reverse engineering problem is feasible, then systematically mapping the possible complex outcomes from the set-ups (the easy problem) might be useful. However, as this paper argues, this can not be assumed, thus the burden of justification in this case is to show that the reverse engineering problem is feasible for this case of system.
- Clearly models of observed complex systems have the potential to be directly relevant to the hard complexity problem. In such work, where the specification of the model set-up comes from evidence about the processes in the observed system, the reverse engineering problem is restricted to only those aspects of the process that are unknown. The more the model is constrained by evidence of the observed system, the easier the determination of questions about the

processes using simulation and other models is. Clearly in this case, the burden of justification is in the extent to which the model is determined by evidence of the observed system and to what extent it is based on assumptions and guesses as to the nature of the processes.

- Papers that don't directly attack the understanding of observed systems but provide tools or critiques that aid one of the above cases is indirectly helpful. I am certainly not arguing that every paper should have to *directly* address the hard complexity problem, merely that there should be some grounds for believing it might *eventually* help in doing so, however indirect that help is. Authors of methodological, philosophical, tool provision, ... etc. papers have to justify that this might be the case.

Clearly with any set of criteria there will be exceptions. There may well be justified exceptions to any established norm of science, however good as a general principle. However in these cases there can be made a special argument why they should be considered, and it does not abnegate a general shift away from papers that only address the easy problem towards ones that address the hard problem. I invite people reading this paper to see how many papers in this conference would satisfy *any* of the criteria listed above.

Conclusion

To put it bluntly, many papers in complexity science are *almost certainly* ultimately irrelevant to any complex system that we might actually encounter. This is not because the papers are stupid, mistaken or their authors did not care about their relevance, but due (I think) to a severe underestimate of the difficulty of the reverse engineering problem. It *was* important to show examples of results that showed how the *easy problem* could be solved in different ways in the earlier period of complexity science, partly because it was novel and partly because it went against some assumptions of science that restricted itself to analytically solvable (i.e. simple) systems. However this has now been comprehensively done. If the field fails to make substantive progress with the *hard complexity* problem then it will become another academic field that gradually loses academic respect, public interest and ultimately its funding. Let us move on to this substantial challenge and address how we might understand and untangle the complexity of observed systems and not just artificial “toy” examples.

Acknowledgements

Some of this work is done under grant XXXX from the EPSRC. My thanks to participants of the EPSRC NANIA project and members of the Centre for Policy Modelling for many interesting discussions.

References

Bedau, weak emergence

Feyerabend, Against Method

Cutland (1980) Computability