Pragmatic Holism[1]

Bruce Edmonds
(b.edmonds@mmu.ac.uk),
Centre for Policy Modelling,
Manchester Metropolitan University,
Aytoun Building, Aytoun Street,
Manchester, M1 3GH, UK.

In the face of complexity, even an in-principle reductionist may be at the same time a pragmatic holist.[2]

Abstract

The reductionist/holist debate seems an impoverished one, with many participants appearing to adopt a position first and constructing rationalisations second. Here I propose an intermediate position of pragmatic holism, that irrespective of whether all natural systems are theoretically reducible, for many systems it is completely impractical to attempt such a reduction, also that regardless if whether irreducible ‘wholes’ exist, it is vain to try and prove this in absolute terms. This position thus illuminates the debate along new pragmatic lines, and refocusses attention on the underlying heuristics of learning about the natural world.

1 Introduction

It appears to me that the reductionist/holist debate is a poor debate. Many of the participants appear not to be seeking truth or useful models about modelling or understanding phenomena but are solely concerned with supporting previously decided positions in the matter, leaving any search for truth solely within their chosen paradigm. The two camps have adopted distinct languages, styles, journals, conferences and criteria for success and thus are largely self-reinforcing and mutually exclusive.

Here I will argue that the concentration on such dogmatic positions centred around largely abstract arguments is unproductive and fairly irrelevant to practical enquiry. In this way I hope to play a small part in refocussing the debate in more productive directions.

I will start by reviewing some of the features of the debate, the versions of reductionism (Section 2.1), some weaknesses in the two sides which make it unlikely that there will be a resolution to the abstract debate (Section 2.2 and Section 2.3) and some irrelevances to it (Section 2.4). I briefly look at some of the general practical limitations to modelling (Section 3) before introducing an illuminating analogy between ordinals and complexity (Section 4). I will argue that the usual definition of computability is too strong (Section 5). Throughout all of the above we see the abstract questions of reducibility coming back down to pragmatic questions which leads me to reject the

1. This article could equally be called “Pragmatic Reductionism”.
2. An adaption of a quote from Herbert Simon: “*In the face of complexity, an in-principle reductionist may be at the same time a pragmatic holist.*” [18].

extreme positions for a more pragmatic approach (Section 6) which will hopefully open up more important and productive questions asked in the conclusion (Section 7).

2 The reductionist/holist debate

2.1 Versions of reductionism

The scientific method is not a well defined one, but one that has arisen historically in the pursuit of scientific truth[3]. From this practice some philosophers have abstracted or espoused a "purer" form of ideal scientific practice, which is epitomized in the reductionist approach. It is around this that debate has largely centred. There are many formalisations of reductionism. Here are some examples:

- "*Any phenomenon can be arbitrarily well approximated by an explanation in terms of microscopic physical laws*"
- "*Every definable process is computable*" (*)
- "*Every causal process is syntactically formalisable*"
- "*Every problem is effectively decomposable into sub-problems*"
- *"The explanation of the whole in terms of its parts"*

All of these are subtly different. They all epitomise a single *style* of inquiry, that any phenomenon, however complex it appears, can be accurately modelled in terms of more basic formal laws. Thus they are rooted in an approach to discovering accurate models of the natural world, namely by searching for simple underlying laws. They range from the abstract question of whether all real systems can be modelled in a purely formal way to more practical issues about the sort of reduction preformed in actual scientific enquiry.

In this paper I aim to show the irrelevance of the abstract question; that when faced with a choice of action it is a very similar range of issues that face both the in-principle reductionist and holist. So for the purposes of this paper I will take the abstract definition (*) as my target absolute definition of reductionism (and hence by implication holism).

2.2 Weaknesses in the reductionist position

Foremost in the weaknesses of the reductionist position is that the abstract reductionist thesis itself is neither scientifically testable nor easily reducible to other simpler problems. Thus, although many scientists take it as given, the question of its truth falls squarely *outside* the domain of traditional science and hence reductionism. Its strength comes from the observation that much successful science has come from scientists that hold this view - it is thus a sort of inductive confirmation. Such inductive support weakens as you move further from the domain in which the induction was drawn. This certainly seems true when applied to various "soft" sciences like economics, where it is spectacularly less successful. The current focusing on "complex systems", is another such possible step away from the thesis' inductive roots.

A second, but unconnected support comes from the Church-Turing thesis. Here the strength of this thesis within mathematics is projected onto physical processes, since any mathematical model of that process we care to posit is amenable to that thesis. If you conflate reality with your model of it then the thesis appears reasonable, but otherwise not.

Thirdly, attempts to formalise any actual scientific reduction in set-theoretic or logical terms, have proved unsatisfactory (see [16]).

3. Others have also arisen, but have been less obviously successful.

2.3 Weaknesses in the holist position

Holist literature abounds with counter-examples to the reductionist thesis. Some of these are seriously intended as *absolute* counter-examples. They tend to fall into two categories: the practical (and so unproved in an absolute sense) and the theoretical but flawed. An example of the former is Rosen's example of protein folding which he justifies as a counter-example to the Church-Turing Thesis (CTT) on the grounds that

> *"...thirty years of costly experience with this strategy has produced no evidence of this kind of stability... despite a great deal of work...the problem is still pretty much where it was in 1960...this is worse than being unsuccessful; it is simply contrary to experience...".* [14]

This is a perfectly valid pragmatic observation, justifying the search for alternative approaches to the subject. It does not, of course, *disprove* the CTT and in itself supplies only weak support for extrapolations to broader classes (e.g. all living organisms).

An example of the latter is Fishler and Firschein [5], where they give the spaggetti, the string, the bubble and the rubber-band computer as examples of machines that "go beyond" the Turing machine. These examples follow a section on the Busy-Beaver problem, which is interpreted as being a function that grows too fast for any mechanistic computation. The examples themselves do not compute anything a Turing Machine can not, but merely exploit some parallelism in the mechanism to do it faster[4]. The implication is that since these "compute" these specific problems faster than a single Turing machine, this is sufficient to break the bounders of the Busy-Beaver problem. Of course, the speed up in these example (which are of a finite polynomial nature) is not sufficient to overcome the busy-beaver limitation, which would require a qualitatively bigger speed-up.

Another example is that used by Kampis [8], that humans can transcend the Goedelian limitations on suitably expressive formal systems. He argues that because any such formal system will include statements that are unprovable by that system but which an exterior system can see are true, and humans can transcend this system and see this, that they thus escape this limitation. He then site's Church's example of the conjuction of all unprovable statements as one *we* can see is true but that is beyond any formal system.

The trouble with this is the assumption that humans can transcend *any* formal system to see that the respectively constructed Goedelian statement is true. Although us humans are quite good at this, the *assumption* that we can amounts to a denial of the CTT already, so this can not possibly used as a convincing counter-example! If you state that the truth of the above is evident to us from viewing the general outline of Goedel's proof, i.e. from a meta-logical perspective, then there will be other unprovable statements from within this meta-perspective. Here we are in no better position than the appropriate meta-logic for deciding this (without again *assuming* we are better and begging the question again).

Church's conjunction of unprovable statements gets us no further. We can only be certain of its truth as an reified entity in a very abstract logic (which itself would then have further unincluded unprovable statements at this level) - otherwise we are merely *inducing* that it would be true based

4. If you take into account the preparation of these devices the speed-up is considerably reduced.

on each finite example, despite that fact that such a trans-infinite[5] conjunction is qualitatively different from these (and undefinable in any of the logics that were summed over).

There are many such examples. To deal with each one here would take too long and distract from the purpose of this paper. Suffice to say that all of these (that I have seen) seem flawed if intended as an *absolute* counter-example to the Church-Turing Thesis.

The basic trouble that the holist faces in arguing against reductionism, is that any argument is necessarily an abstraction. This abstraction is to different degrees formal or otherwise. To the degree that it is informal it allows equivocation and will not convince a skeptic. To the degree that it is absolute/formal it comes into the domain of mathematics and logic where the Church-Turing thesis is very strong (by being almost tautologous). While informal arguments can be used with other holists, in order to argue with a reductionist a more formal argument seems to be required[6].

It appears that it is a necessity limitation regarding the nature of *expression itself* that makes any such *complete* demonstration impossible.

2.4 Irrelevances to the debate

Associated with the debate on the absolute question, but not central to it, are a host of old-chestnuts that have not been shown to be relevant, but are often assumed to be crucial. There is not space to deal with them all or thoroughly enough to convince a believer of their inadequacy, but I list some of the more frequent of them below. Despite their weakness to determine the *absolute* question, they each have strong *practical* consequences.

2.4.1 Determinism

Whether natural systems are deterministic or not, in an absolute sense, seems to be an untestable question. Both the deterministic and indeterministic viewpoints adequately describe the observed world. Thus the abstract question of whether a real system is deterministic must be irrelevant to the abstract question of the validity of the reductionist thesis[7].

Artificial situations can be categorised as deterministic or otherwise. For example in a game your next move may be determined by the rules or you may have a choice. Within the framework you are considering, there is either a mechanism for determining your move (i.e. the player's strategy for the game) or not. If there is, then the move is determined by that, if not, it is undetermined - i.e. there is simply not enough information to determine this by any process (mechanical or otherwise). Note that whether the move is determined *does* depend on the framework you are considering, but within any particular framework (however general), whether something is determined is not relevant to the *absolute* question of whether the situation is reducible or not.

However, this does highlight the practical importance of choosing the appropriate framework for a problem. The framework greatly effects the practicality of modelling a system. In the above example in a framework which includes the players strategies, it may be possible to model the game but impractical if you choose merely the rules and possible sequences of moves.

2.4.2 Analogue vs. digital

There is a basic difference between what can be theoretically modelled using analogue and formalisms[8]. Sometimes this is on the grounds of the importance of noise (for this see Section 2.4.4 below).

5. Since "formal systems" would include models with sizes of all the hierarchy of ordinals, the size of the conjuction would be strictly bigger than any infinite ordinal!
6. Despite the fact that this does not seem to be required of the reductionist stance by reductionists!
7. Unless this question is also as untestable and irrelevant as that of determinism.
8. For example [17], even here any practical computation whose resutl is measurable by us would not transcend a Turing Machine.

There *is* an essential difference between analogue and digital in the abstract. You can not encode all analogue values as digital, only some of them. This can be proved with a classical diagonal argument. This means that literally we can not talk about most analogue values, except as a collective abstraction (*"let x be a real number..."*), as there are no finite descriptions of them.

The digital and analogue can arbitrarily approximate each other, thus the colour of a pixel on a VDU is composed of different wavelengths (analogue), which is encoded by the computer as a binary number (digital), which is encoded as voltages in circuits (analogue), which correspond to energy levels (digital).

The natural world may, at root, be analogue or digital, we do not know[9]. Even with matter and energy one could argue that the quanta are a result of observing a continuous wave function. Thus arguments which rely on a fundamental difference between the analogue nature of reality and the digital nature of formalisms and the modern computer, must be somewhat arbitrary.

This is not to say that either simulating the analogue by the digital (or visca versa) does not present considerable practical difficulties.

2.4.3 Ability to modify hardware

The ability of an organism to modify its own "hardware", for example when a protein acts on its own DNA (e.g. to repair it), or at least acts to effect the interpretation of that DNA into proteins, is sometimes compared to a Turing machine which cannot directly effect its "hardware" (as usually defined).

This separation of hardware and software is arbitrary unless "physicality" can be shown to be important attribute, effecting what can be computed. Otherwise, there is nothing to stop a Turing machine simulating such a change in hardware (including its own). For example, imagine a Turing machine which could execute an instruction-type that could change one of its own instructions. It would seem at first glance that this new machine "goes beyond" the usual version, but this is not so. A normal Turing machine can compute exactly the same functions as the new enhanced machine, because although it cannot change its own instructions, those simple instructions can be combined in a sophisticated way to simulate the computation of the enhanced machine.

Several such "essential" characteristics of such physicality are possible.

1. The presence of noise in analogue systems (for this see Section 2.4.4 below).
2. A fundamental difference between matter and symbols (as in [10]). This is closely connected with the problem of measurement.
3. That arbitrarily small changes in the initial conditions have significant effects on the outcomes. The significance of this is either that noise (for this see Section 2.4.4) can then be significant or that due to its analogue nature you can never know the initial conditions sufficiently (for this see Section 2.4.2 above).

We are thus left with the argument as to a fundamental dichotomy between matter and symbols. Whether or not this turns out to be a fundamental distinction, it is not clear why this would effect of the ability to modify one's own hardware (or simulate such a modification in software).

2.4.4 Noise and randomness

Noise is a random input into a system's processes. Such randomness can be defined in several ways. It can be any sufficiently variable data which originates from outside the scope of a system's model of its world, and is thus unpredictable. It can be data which passes a series of statistical tests. It can be a pattern which is incompressible by a Turing machine.

9. Of course, it may be neither.

In any case there are fully deterministic processes which produce sequences that are practically indistinguishable from random ones from any particular system's point of view. Thus a system with noise can be simulated by a model with the addition of such a process, such that the system does not have full access to the workings of that process.

In the opposite direction a noiseless system can be arbitrarily approximated by one that has noise, by suitable redundancy in its construction. This is how we maintain information in digital computers and our own genome.

2.4.5 Particular formal languages (2-valued logic)

One particular bugbear of holists, is classical two-valued logic. This is criticised as not being expressive enough to capture all the meaning orreasoning necessary to model some systems. Sometimes it is Zermelo-Frankel set-theory that is the target. For example Rosen [14] suggests a approach to modelling in terms of category theory as a possible way forward in modelling complex biological systems.

Often it seems that the importance of *whether* a formal system is applicable is based on a shallow reading of the formal system's immediate properties and passes-over what further expressive features can be formalised within it.

In fact the choice of formal system is not critical in absolute terms, as long as the system is expressive enough. For example category theory and set theory can both be used to formalise the other [12], similarly Classical first order logic can be used to formalise almost any other logic [6]. So there is no fundamental *absolute* grounds for preferring one such formal system to an other[10].

2.4.6 Self-reference

Another way in which holists claim that there are systems that are absolutely unamenable to formalisation is by exhibiting those that involve some form of self-reference. An example of this is found in [9].

There are two different forms of self-reference *total* and *grounded*. Total self-reference is completely self defining, if you follow the causal (or formal) chain backwards, you do not come to a fixed atomic starting place, but find an infinite recursion of definition in terms of itself. Grounded self-reference starts at a specific place (i.e. state or set of axioms), and then the next state is defined in terms of the last state etc. so after a while the current state is almost completely defined in terms of itself and the origin is lost for all practical purposes.

Most examples of self-reference in the natural world would seem to be cases of grounded self-reference: life itself presumably started from some point, which arose from non-living state[11]; language is ultimately grounded in our shared experience, either directly as a child learns its first language, or indirectly in the evolution of language in our species[12]; even the universe itself seems to have passed through an initial equilibrial stage [21].

It would seem that total self-reference is difficult to embed in a traditional formalisation[13], if only because in formalising something you need somewhere to start from. It is possible to dissect such a system so that it is representable *within* a traditional framework (e.g. the technique described in [9]), but only by effectively grounding it. Thus even if total self-referential formal systems exist they are only usable in modelling and easily communicable if grounded.

---

10. There will be very strong pragmatic grounds for the choice of system, like tractability or a wish for a *natural* interpretation.
11. If one considers "life" as one capable of varying degrees then this argument still holds, suitably amended.
12. Such a mechanism as Harnad suggests in [7].
13. This is different from the consistency of self-reference, this is relatively easy to establish, as in [10].

Grounded self-reference is formalisable by traditional formal systems, even though in some cases this may be a cumbersome and "unnatural" way to proceed. It is true that some such formalisations are either resistant, or do no have, analytic solutions that allow effective prediction of future (or even description of past) behaviour, but this has always been true of even the most classical of formalisms and thus is not relevant to the *absolute* reductionist/holist question.

Again whether we use traditional styled or (grounded) self-referential systems for modelling, relates not to abstract but pragmatic considerations.

2.4.7 Simultaneity

In most existing computation devices, computation proceeds sequentially. Even parallel devices are usually arranged so that their computations are equivalent to such a sequential approach. Likewise in almost all formal systems, facts are derived via an essentially sequential proof. Even when the proof is not sequential in nature, its verification is.

Natural systems, however, seem to work in parallel. Von Foerster gives an example of a box with many block magnets in it[14], the box is shaken and when opened they are arranged in a very non-random way, resulting in an attractive sculpture to an exterior observer. The two views of the box, internal and external are simultaneous and different. It is claimed that such simultaneous and (in some cases) irreconcilable viewpoints, mean that a single consistent formalism of a meta-model incorporating both viewpoints is impossible.

If you have a parallel system there will be either some conflict avoidance or a conflict resolution mechanism (where by "conflicting" I mean exclusive). Of course, it is quite possible to have cases where (as in the above box of magnets example) there are views that appear to be conflicting, but you won't have conflicts within the same context, this is impossible if a consistent language is used. Here it is not the simultaneity that is the problem but the reconciliation (or lack of it) of the same thing from within different frameworks (see Section 3.4 below).

On the other hand, the difficulties of reconciling different simulataneous streams means that often the only practical option is to accept such different views as complementary.

3 Practical limits to modelling

Although it may be hard to prove practical limits to modelling any specific problem, there are many general practical limitations.

3.1 Finiteness

It seems we (us and our tools) are part of a finite universe, and are thus also finite. Any model we make, use or understand will also be finite[15]. Quite apart from this our formal communications (written articles) are definitely finite. Thus any *practically* useful model that we want to share will also be finite.

In these circumstances the fact that an Turing machine (which is *essentially* infinite) could compute something, may not be relevant if the mapping from this abstraction to an actual computer may mean that the computation is impractical. Thus the abstract question of the CTT is superseded by the question of the practicalities of modelling.

3.2 Limited computational resources

As well as limited memory we also have a limited time to do the computations in. It has been calculate that quantum mechanics imposes a limit of 10^{47} bits/gram/sec on the amount of

14. From a conference paper called "self-Organising Systems and their Environments" in 1960.
15. Although it may well be most useful to represent or talk about an abstraction of this as infinite.

information that can be computed by each gram of matter per second [2]. Taking a conservative guess at the total mass in the universe as 10^{70} grams and the total time before the heat death of the universe as 10^{15} years this gives us an upper limit of about 3×10^{142} total bits of computation that could possibly occur in this universe. This would be insufficient to even investigate the possible colourings of a 12 by 12 checker-board using just 10 colours.

Thus problems which take undue computational time come up against a fairly fundamental computational limit, even if they are *theoretically* computable.

3.3 Complexity

Computational Complexity is concerned with the computational resources required, once the program is provided. It does not take into account the difficulty of writing the program in the first place. Experience leads me to believe that frequently it is the writing of this program that is the more difficult step.

More fundamental is what I would call “analytic complexity”. This is the difficulty of analysing (producing a top-down model) of something, given a synthetic (bottom-up) model [4]. Whether or not this difficulty is sometimes ultimate, few people would deny that such difficulties exist and that the exist arbitrary levels. Given that our analytic capabilities will always be limited (see above in Section 3.2), such complexity will always be a practical barrier to us[16].

3.4 Context

Not all truth can be expressed in a form irrespective of context. The very identity of some things (e.g. society) are inextricably linked to context. Thus we will have to be satisfied that, for at least some truths, it will not be practical to try and express them in a very general context and hence acquire the ‘hardness’ of more “analytic” truths (like “all bachelors are men”). It is true that we can laboriously express larger and larger meta-contexts encompassing sub-contexts, but this will involve the construction of more and more expressive languages[17] and require disproportionately more computational power - this will make this sort of endeavour impractical, beyond a certain level[18]. Choosing an appropriately restricted context is one of the most powerful means at our disposal for coping with otherwise intractable situations.

4 The number - complexity analogy

Rosen introduces an analogy between what he calls complexity (i.e. things that aren’t mechanisms) and infinity; the reductionist/syntactic approaches to modelling correspond to finite steps. He claims that many systems (including all living organisms) are unameanable to such steps and qualitatively different - they correspond to infinity. Thus he postulates that to model these “complex” systems require some transcendental device, like taking limits or some form of self-reference.

I wish to alter this analogy and hopefully deepen it. I wish to take an analogy between numbers and complexity. This size corresponds to the difficulty of modelling a system in a descriptive top-down fashion given a language of representation and almost complete information (model) from the bottom-up perspective of it components[19]. Thus infinite size would correspond to infinite such difficulty - i.e. impossibility of such modelling (which roughly corresponds to Rosen’s

16. For a classic account of the inevitable emergence of complexity in biological systems see [22]
17. For example the logical structure in [1].
18. For an overview on this problem in biology see [23].
19. For more on the definition of complexity see [4].

"complexity"). The abstract debate would then correspond to the question *"Are there systems with infinite complexity?"*.

Here we need to examine what we mean by the *existence* of such systems. The problems of showing that such systems exist are remarkably close to those involved in showing that infinity *exists*. You can not exhibit any real manifestation of infinity, since the process of exhibiting is essentially finite. Even if we lived in a universe that was infinite in some respect, you could not show a *complete* aspect that was infinite, only either that an aspect appeared *unbounded* or that a reasonable projected abstraction of some aspect was infinite.

Note that I am not saying that infinity is *meaningless*, merely that it is always an abstraction of reality and not a direct exhibitable property of any thing. That infinity is a very useful abstraction is undeniable - it may be possible to formulate much of usable mathematics without it, but this would surely make such symbolic systems much more cumbersome. So when we say something is infinite, we are talking about an abstract projected property of our model of the item, *even if the thing is, in fact, infinite*. It is just that *exhibiting* is essentially a finite process.

I suspect that the same is true of the irreducibly complex. A language of irreducible "wholes" is useful in the same sense that infinity is useful, but *only* as an abstraction of our model, *irrespective of whether these "wholes" exist*. If they do not exist, the language of the holist is still useful as an abstract shorthand for systems whose complexity is potentially unbounded. If they do exist the language of "wholes" would still be necessarily abstract, i.e. *not* referring to direct properties of real things, *even if the systems referred to were irreducible*. It is just that *exhibiting* such systems (especially formally) is essentially a reductive process.

5 Formal modelling and reductionism

The computability of a well-defined function is a purely formal question. A function is computable if an index *exists* such that a universal Turing machine (UTM) with that index calculates that function. For example, consider an enumeration of halting computable programs in order of size. Every function represented by programs in this class is computable without us being able to compute (or "write") the programs for these functions, otherwise we could use this enumeration to compute the "halting problem" [19][20].

This shows that just because if some function is theoretically computable does not mean that we have the means to find the program to compute it. In other words, the characterisation of reducible as computability is *too strong* for actual use in reducing a problem. To be able to *really compute* something we have to be able to follow the instructions (or have a machine do it) *and* write the program to do this (or *really compute* that). This can form a very long chain (the program that computes the program that computes the program that...) , but eventually it has to be grounded in a program we are able to write ourselves, *if the final program is to carry out our intentions*. I call this "intensional modelling" (or "intensional computing", depending on the context). This is what *we* can compute (using computers as a tool), as opposed to what *can theoretically* be computed by a Turing machine.

If we take a pragmatic view of reductionism only as such *intensional* modelling, then we come to the surprising conclusion that there may be some computable functions that *we* can't compute (intentionally), *even by using a computer.* We may, of course, come across a few of these programs *accidentally*, for example by genetic programming, but we can never be certain of this and verifying that these programs meet a specification is itself uncomputable in general (although you may be able to analyse a few such particular examples sufficiently).

20. See [3] for a more accessible account.

This raise the possibility that if a form of self-replicating, evolving software life appeared by accident (say as the unintentional result of a genetic programming program) this may be as difficult to model and understand as its more tangible cousins. The code may be so complex and self-referential that it was *as* difficult to decode as the mechanisms in normal life (as we know it). So even if (a suitable version of) the reductionist thesis were true in the abstract, we might still be forced to use more holistic or uncertain methods to model the phenomena we come across.

6 Heuristics in the search for truth

So whichever is our belief about the abstract reductionist/holist question, we are left with very similar pragmatic choices of action when faced with an overly complex problem. Here reductionist techniques will be of little practical value for *us* as limited beings and we have to look to other alternatives if we want to make progress on them. Whether you choose another (possibly less successful) approach, depends upon the trade-off between the difficulty of reduction and the importance of progress (of what ever kind) being made onthat problem. In the end, the biggest practical difference between a reductionist and a holist is often only that a reductionist then chooses another problem where the reductionist technique has more chance of success and the holist chooses alternative avenues of attack upon the same problem.

The point is that there is no necessity to prejudge this decision for every case, neither to always say that alternative types of knowledge are worthless, despite the importance of the problem nor to say that it is never worth abandoning a problem because of the type of knowledge that is likely to be gained about it. I call these the *extreme reductionist* and *extreme holist* positions respectively.

To hold to the extreme reductionist position in a practical sense, one must surely claim that no problem is so much more important than other more susceptible problems to be worth swapping the sort of analytic knowledge that results from reductionist approaches for other types of knowledge. This can be a result of one of several subsidiary claims:

1. that all problems are practically susceptible - this would amount to denying any practical limitations upon ourselves at all;
2. that there are always a near indefinite supply of equally important problems - denying any real difference in the importance of problems, regardless of circumstance;
3. or that alternative forms of knowledge are always effectively worthless - presumably including the reductionist thesis itself!

To hold to the extreme holist position in a practical sense, one would have to claim that either there was no advantage to reductionist knowledge as compared to another type in any circumstances or that a problem was *so* important that it was not appropriate for anyone to research any other, more amenable, problems.

These would both be extreme positions indeed! I know of no one that holds them in these forms. The rest of us fall somewhere in between *in practice*: we accept that there are some worthwhile problems where the reductionist technique works well and we also accept that there are problem domains where the chances of a reductionist technique working are so remote and the problem so important that we would value other forms of knowledge about it.

This does not mean that we will all have the same priorities in particular cases, just that these decisions are essentially a *pragmatic* ones differing in degree only. Once attention switches from the sterile abstract question of whether *in principle* all problems are amenable to a reductionist approach (and thus implicitly excluding the extreme positions outlined above), we can start to consider the rich set of possible strategies for making such choices in different cases[21]. This is has been up to now a largely uncharted area, but one that might pay rich dividends.

7 Conclusion - combining a plurality of techniques

The conclusion is thus, maybe disappointing: you need to think about the applicability of the reductionist technique before using it (just like any other technique) *regardless of the answer to the abstract reductionist/holist question*.

Such an awareness opens up the possibility of a more systematic study of ways of combining techniques for the successful elicitation of knowledge. Ways which would include developing more effective ways of using different types of knowledge to further guide and refine that search. Such practical heuristics are maybe all that we finally have - rejecting both blinkered single-strategy approaches and the extreme relativism of *anything goes*.

Acknowledgements

I would like to thank the members of the Principia Cybernetica mailing list[22] with whom I have had many useful discussions on this subject, especially Don Mikulecky, Jeff Prideaux, Cliff Josyln and Francis Heylighen.

References

[1] Barwise,J; Perry,J (1983): Situations and Attitudes. MIT Press, Cambridge, MA. 352 pages.

[2] Bremmerman,HJ (1967): Quantal noise and information. 5th Berkeley Symposium on Mathematical Statistics and Probability, 4: 15-20.

[3] Cutland,NJ (1980): Computability. Cambridge University Press, Cambridge.

[4] Edmonds,B (forthcoming): What is Complexity? - the philosophy of complexity per se with application to some examples in evolution. In Heylighen,F; Aerts,D (eds), The Evolution of Complexity, Kluwer, Dordrecht. Also available electronically at **`http://www.fmb.mmu.ac.uk/~bruce/evolcomp`**.

[5] Fishler,MA; Firschein,O (1987): Intelligence in eye, the brain and the computer. Adison-Wesley, Reading, MA.

[6] Gabbay,DM (1994): Classical vs Non-classical Logics (The Univiversality of Classical Logic). In: Handbook of Logic, in Artificial Intelligence and Logic Programming. Vol. 2. (Eds: Gabbay,DM; Hogger,CJ; Robinson,JA), OUP, Oxford, 359-500.

[7] Harnad,S (1990): The symbol grounding problem. Physica D 42, 335-346.

[8] Kampis,G (1991): Self-Modifying Systems in Biology and Cognitive Science: A New Framework for Dynamics, Information and Complexity. Pergamon Press, Oxford. 565 pages.

[9] Kampis,G (1995): Computability, Self-Reference, and Self-Amendment. Commuincation and Cognition - Artificial Intelligence 12, 91-109.

[10] Lofgren,L (1968): An Axiomatic Explanation of complete Self-Reproduction. Bull. Math. Biophys. 30, 415-425.

[11] Pattee,HH (1995): Evolving self-reference: matter, symbols and semantic closure. Communication and Cognition - Artificial Intelligence, 12, 9-27.

[12] Marquis,J-P (1995): Category Theory and the Foundations of Mathematics: Philosophical Excavations. Synthese 103, 421-447.

[13] Rosen,R (1985): Organisms as Causal Systems Which Are Not Mechanisms: An Essay into the Nature of Complexity. In: Theoretical Biology and Complexity: Three Essays on the Natural Philosophy of Complex Systems. (Ed: Rosen,R) Academic Press, London, 165-203.

21. Sarkar [16] has even suggested that there are several different types of models of reductionism and so our model of reductionism itself must be chosen pragmatically.
22. PRNCYB-L, for details see: http://pespmc1.vub.ac.be/MAIL.html

[14] Rosen,R (1991): Life Itself - A Comprehensive Enquiry into the Nature, Origin and Fabrication of Life. Columbia University Press, New York.
[15] Rosen,R (1993): Bionics Revisited. In: The Machine as Metaphor and Tool. (Eds: Haken,H; Karlquist,A; Svedin,U) Springer-Verlag, Berlin, 87-100.
[16] Sarkar,S (1992): Models of Reduction and Categories of Reductionism. Synthese 91, 167-194.
[17] Siegelmann,HT (1995): Computation Beyond the Turing Limit. Science 268, 545-548.
[18] Simon,HA; (1981), The Architecture of Complexity, in The Sciences of the Artificial, Eds., MIT Press, Cambridge, Massachusetts, pages 192-229.
[19] Turing,AM (1936): On Computable Numbers, with an application to the Entscheidungsproblem. Proc. London math. Soc., 2 42, 230-265.
[20] Von Neumann,J (1966): Theory of Self-Reproducing Automata. University of Illinois Press, Urbana, Illinois.
[21] Weinberg,S (1977): The First Three Minutes. Basic Books, New York.
[22] Wimsatt,W (1972): Complexity and Organisation. In: Studies in the Philosophy of Sciences. Vol. 20. (Eds: Schaffner,K; Cohen,R) Reidel, Dordretch, 67-86.
[23] Yates,FE (1978): Complexity and the limits to knowledge. American Journal of Physiology 235, R201-204.