

TYPES OF RESEARCH in COMPUTING SCIENCE SOFTWARE ENGINEERING AND ARTIFICIAL INTELLIGENCE

**Based on a conference organised by CPHC
at Manchester University in January 2001.**

Maintained by Aaron Sloman¹

<http://www.cs.bham.ac.uk/~axs/>

Last updated: 18 Dec 2007

Derived from: <http://www.cs.bham.ac.uk/~axs/misc/cs-research.html>

Contents

1	Background (and Updates)	2
2	The Five Categories of Research in CS and AI	4
3	Comments on the categories	5
3.1	The study of what is possible – and its scope and limits	5
3.2	The study of existing (naturally occurring) information-processing systems . . .	5
3.3	Research involving creation of new useful information-processing systems. . . .	7
3.4	Creation and evaluation of tools, formalisms and techniques to support these ac- tivities	8
3.5	Research on social and economic issues	8
4	Evaluation Criteria	9
4.1	The study of what is possible – and its scope and limits	9
4.2	The study of existing information-processing systems	10
4.3	Research involving creation of new useful information-processing systems. . . .	10
4.4	The creation of tools, formalisms and techniques	11
4.5	Research on social and economic issues	11

¹Please send comments, criticisms, suggestions, corrections, etc. to
Aaron Sloman, School of Computer Science, The University of Birmingham, B15 2TT, UK
EMAIL A.Sloman@cs.bham.ac.uk

1 Background (and Updates)

A Research Strategy conference was organised by the CPHC (Conference of Professors and Heads of Computer Science) at the University of Manchester 6-7 Jan 2000. It was attended by about 100(?) people (not only professors and heads). On the first afternoon there was an introductory panel session concerned with how the Computing Science community should present its research objectives and achievements to EPSRC and the bodies which award funding to EPSRC.

During the ensuing discussion I suggested a high level way of dividing up research aims into four main categories (later expanded to five), which, in part, need to be evaluated differently.

Both during the conference and subsequently I received comments and requests for clarification and references. So I thought I should write down what I said (as far as I remember it), expand it a bit, and circulate it for comment and criticism.

The resulting document is in this file at <http://www.cs.bham.ac.uk/~axs/misc/cs-research.html>² Whenever I have afterthoughts, or receive criticisms, comments and suggestions for improvements, I may modify/correct/extend the file, with acknowledgements where appropriate in footnotes.

NOTE Added 18 Dec 2007

A paper written by Allen Newell addresses some of the issues listed here:

A. Newell (1983) Intellectual issues in the history of artificial intelligence, in *The study of information: interdisciplinary messages*, pp. 187–227, Eds. F. Machlup and U. Mansfield, John Wiley & Sons, New York,

Available in the Newell Archives

<http://diva.library.cmu.edu/webapp/newell/item.jsp?q=box00034/fld02334/bdl0002/doc0001/>

(Nearly 25 years later, it seems that the

NOTE Added 2 Nov 2007

Alan Bundy has developed a web site which serves some of the same purposes as this one:

<http://www.inf.ed.ac.uk/teaching/courses/irm/notes/hypotheses.html>

The Need for Hypotheses in Informatics

NOTE Added 28 Feb 2006

Most of this document was written before the UKCRC initiative on Research Grand challenges: <http://www.ukcrc.org.uk/grand.challenges/>.

Several of the grand challenge proposals that emerged within that initiative are examples of the proposals made in this document to think of the research problems as much broader than traditional computer science, especially

- GC1: In Vivo – In Silico
See URL <http://www.cmp.uea.ac.uk/Research/ivis/index.jsp>
- GC5: The Architecture of Brain & Mind
See URL <http://www.cs.bham.ac.uk/research/projects/cogaff/gc/>
- GC7: Journeys in Nonclassical Computation
See URL <http://www.cs.york.ac.uk/nature/gc7/>

NOTE Added 19 Dec 2004

²The ideas here overlap with those that went into the overview of AI which was produced (with help from colleagues in various places) for the QAA computing science benchmarking panel: <http://www.cs.bham.ac.uk/~axs/courses/ai.html>

The UKCRC Grand Challenge initiative, proposed in 2003 and described in <http://www.nesc.ac.uk/esi/events/Grand.Challenges/proposals/> illustrates some of the points made below about different kinds of research.

Discussions of Grand Challenge 5 ('Architecture of Brain and Mind' <http://www.cs.bham.ac.uk/research/cogaff/gc/>), which is one of the long term grand challenges with no definite end point (like many scientific and medical grand challenges), raised the difficult question of how to identify progress.

This is an issue addressed in a very relevant way in the writings of Imre Lakatos, who extended some of the ideas of Karl Popper by making a distinction between 'progressive' and 'degenerating' research programmes, where the important point is that it may be *impossible* to decide whether a research programme is of one type or another at early stages in the programme: the decision requires analysis of an extended period of research. There are many internet sites discussing, summarising, criticising or reproducing Lakatos papers.

I. Lakatos (1980),
The methodology of scientific research programmes, in
Philosophical papers, Vol I, Eds. J. Worrall and G. Currie,
Cambridge University Press,

See

<http://www.personalityresearch.org/metatheory/lakatos/progressive.html>

<http://www.stephenjaygould.org/ctrl/lakatos.prediction.html>

<http://www.answers.com/topic/imre-lakatos>

In the context of Grand Challenge 5, I offered a scenario-based methodology that is useful both for planning research and for evaluating it, based on development of a large collection of (partially) ordered scenarios of varying depth and difficulty.³

NOTE Added 9 Jan 2001

Following recent discussions about UK CS research on the cphc-members email list, and a note circulated by Alan Bundy referring to a list of research topics produced some time ago by a CPHC committee.⁴

I have added a new category of research topics, "Research on Social and Economic Issues". So although there were originally four categories, there are now five, although the new one is not a sub-discipline of Computer Science but rather a multi-disciplinary research area.

³The methodology is summarised in <http://www.cs.bham.ac.uk/research/cogaff/gc/targets.html>
It is also being used in connection with an ambitious EU-funded project that began in September 2004, described in <http://www.cs.bham.ac.uk/research/projects/cosy/>

⁴At the time this note was written, the CPHC document on Generic Questions by Alan Bundy was accessible at <http://homepages.inf.ed.ac.uk/bundy/tmp/generic-questions.html> However, it may not remain publicly available.

2 The Five Categories of Research in CS and AI

Research in Computing Science and AI falls into four main categories, with different types of aims, and different success/failure criteria, though each of the categories feeds on and contributes to the others, and there are some kinds of research which straddle categories. There is a fifth cross disciplinary category which is of great interest to many computer scientists though it is not strictly a part of Computer Science or AI, though concepts and techniques from both form part of its subject matter and can also be used to further its aims.

1. **The study of what is possible – and its scope and limits**
Including both mathematical and less formal modes of theorising.⁵
2. **The study of existing (naturally occurring) information-processing systems**
E.g. animals, societies, brains, minds, (Sometimes described as “Natural Computation”).
3. **Research involving creation of new useful information-processing systems.**⁶
I.e. research directly related to engineering applications.
4. **Creation and evaluation of tools, formalisms and techniques to support all these activities.**
5. **Research on social and economic issues**
Including studies of the social and economic impact of computing and AI, ethical issues, changing views of humanity, etc.

These categories are described in more detail below.

Because different kinds of activity need to be evaluated in different ways (see below), there are implications regarding how EPSRC ought to organise its reviewing of grant proposals, and perhaps also implications regarding what proposers should say about their objectives.

In particular, we should strongly resist real or imagined pressures to force all our research into Category 3, and should not be tempted to disguise research in the other categories, or justify it merely as a contribution to Category 3.

Added 30 Mar 2005:

In the light of recent discussions on the CPHC email list it may be worth subdividing Category 3 in various ways. E.g. some of the research contributions to practical applications involved a relatively simple yet new and powerful key idea (e.g. the original idea of the World Wide Web), whereas others are inherently concerned with production of something large and complex requiring the development of a large and complex collection of ideas, e.g. the design of a secure and robust air traffic control system, or a novel nationwide information system for the health service. Many such systems require the use of knowledge and techniques from many disciplines.

The next section explains in more detail what the above categories are, and how they are related and mutually dependent. The section after that explains how the evaluation criteria relevant to these categories of research differ and where they overlap.

⁵I am grateful to Jim Doran for reminding me of the need to allow less mathematical work in this category, especially as it applies to most of my work as a philosopher doing AI!

⁶Michael Kay, ICL, pointed out that my original title for the third kind of research (“Creation of new useful information-processing systems”) was misleading. Many people work on creating new useful information-processing systems but are not doing research. The description in section 3.3 was rephrased to accommodate his comments. Rachel Harrison, Reading University, suggested including evaluation in this category.

(In what follows I use “type” and “category” interchangeably as terms of ordinary English, not as technical terms.)

NB: where lists of examples are given they are merely illustrative and are not intended to be exhaustive, or to define a category.

3 Comments on the categories

3.1 The study of what is possible – and its scope and limits

This includes a lot of work using mathematics and logic, such as work on semantics of computation, and theorems relating to limits of computation, complexity, properties of mechanisms for cryptography, mathematical analysis of different classes of computations, studies of the expressive power of different formalisms, analysis of properties of various kinds of information-processing architectures, network protocols, scheduling algorithms, etc. etc. Much of this work involves the study of types of virtual machines and their properties. They need not be machines which could exist in nature: e.g. some might be infinite machines.

This category also includes less formal, and possibly less rigorous[3] (cf. Section ??), exploratory investigations of new types of architectures, including virtual machine architectures, hardware and software mechanisms, forms of communication, ontologies, etc. in order to investigate their properties and their trade-offs. Examples in AI include explorations of various forms of representation or high level architectures for use in intelligent systems. Sometimes work that starts off in this informal way leads to new formal, mathematical developments, as has happened throughout the history of mathematics.

Often work in Category 1 builds on and abstracts from experience gained in tasks in the other categories, just as much of mathematics derives from attempts to find good ways of modelling complex physical structures and processes, e.g. Newton’s and Leibniz’ invention of Calculus, and the early work on probability theory inspired by gambling devices.

Very often this theoretical work addresses problems that are sufficiently complex to require the use of tools of the sorts developed in research of Type 4.

Purely theoretical work often develops in such a way as to provide concepts, models, theorems and techniques relevant to the other three kinds of research, though even if it does not do so it can still be of great interest and worth doing as a contribution to human knowledge. It has intrinsic value comparable to that of music, poetry, painting, sculpture, literature, mathematics and dare I say philosophy.

3.2 The study of existing (naturally occurring) information-processing systems

(Sometimes described as “Natural Computation”.) This is scientific research of another kind: the attempt to understand, explain or model, things that exist in the world, as opposed to exploring what is possible (Category 1) or finding ways of creating new useful things (Category 3). Of course such understanding can sometimes lead to useful practical applications, by enabling us to predict, control or modify some of the behaviour of systems after we understand them. But that is not a *requirement* for the work to be of great scientific value (though it can be part of the selection process where there are competing theories).

There are many kinds of naturally occurring systems, including machines that manipulate matter, machines that manipulate forces and energy and machines that manipulate information – including virtual machines that cannot be observed and measured as physical machines can.

Long before there were computers or computer science there were many types of extremely sophisticated information-processing systems, including animal brains, insect colonies, animal societies, human social and economic systems, business organisations, etc. More recently new systems have grown which are enabled by information-processing artefacts, but are as much natural systems worthy of study as a society or the weather, for instance traffic systems or the internet.

The processes of biological evolution form another such naturally occurring information-processing system. Over huge timescales, using mechanisms which are still only partially understood, it compiles information about many types of environments and many kinds of tasks (e.g. serving needs of organisms) into a diverse collection of wonderfully complex and extremely successful designs for working systems, far exceeding in complexity, sophistication and amazing robustness, anything yet produced by human designers of information-processing machines.

Some physicists argue that even the physical universe is best construed as ultimately composed of information-processing systems, not yet fully understood. Whether work in computing science will contribute to that understanding I do not know, though there are attempts in that direction.

Prior to the development of computing science the study of complex naturally occurring information-processing systems was often very shallow, mostly just empirical data-collection, often using theories expressed only in crude general forms or coarse-grained equations or statistical correlations which failed to capture or explain any of the intricate detail of processes observed. Since the middle of the last century, the study of different forms computation has enriched our ability to find new ways of formulating and testing powerful models and theories for explaining and predicting natural phenomena.

Information-processing models and theories are being developed in many scientific domains, as people find that they provide richer, more powerful explanatory capabilities than the old paradigms (e.g. equations relating observed of measurable quantities). This in turn is feeding new ideas into computing science.

This has most obviously happened over the last 50 years or so in work in Artificial Intelligence, a discipline whose scientific “arm” has in the past mainly focused on attempts to model and explain aspects of human-intelligence, though there are increasingly attempts at modelling various kinds of animal intelligence. (See the overview of AI in <http://www.cs.bham.ac.uk/~axs/courses/ai.html>⁷.) Unfortunately many psychologists have no appreciation of this as shown by the pressures by which the British Psychological Society causes Psychology departments to stop allowing their students to take AI courses, which are not recognised as relevant. (Behind all that is an out-dated philosophy of science based on an incorrect model of physics as a science that collects lots of measurements and then searches for correlations.)

A more recent development is the growing interest in interpreting biological evolution as a form of information-processing which has also inspired exploration of novel forms of computation which may or may not turn out to be useful for modelling nature.

It is arguable that the activity of engineers, working individually or in teams, is an example of a naturally occurring process and therefore empirical investigations of different kinds of practices, methodologies languages, tools etc., and how they work, could fit into Category 2. This is usually

⁷See URL <http://www.cs.bham.ac.uk/~axs/courses/ai.html>

an intrinsic part of research in Category 4, which is primarily intended to support Category 3. However, analysis and simulation of human engineering activities can fit into Category 2, and work in AI/Cognitive science on simulation of human design processes would clearly do so.⁸

3.3 Research involving creation of new useful information-processing systems.

Research closely related to production, analysis and evaluation of practical applications is the main **engineering** branch of computing science, though Category 4 also includes a type of engineering. Category 3 overlaps with Category 2 insofar creation of explanatory theories and models often involves designing and implementing new and complex systems requiring significant engineering skills. There is also overlap insofar as building useful devices often requires a deep understanding of the environment in which they are to operate. E.g. many software engineering projects producing systems to be used by or interact with humans, including HCI projects, have failed because they used shallow and grossly inadequate models of human cognition, motivation, learning, etc.

Despite some overlap with Categories 1 and 2, the primary goal of research in Category 3 is not to study theoretically possible systems and their properties, nor to help us understand already occurring information-processing systems. The goal is to enable us to create new practically useful systems, which may either:

(a) provide new (or improved) types of artefacts capable of performing functions that were previously performed only by natural systems such as humans and other animals (e.g. doing numerical computations, proving mathematical theorems, translating from one language to another, designing new machines, managing office records, recognising faces)

or, increasingly often,

(b) develop systems to perform tasks that could not be achieved at all previously, e.g. the construction of global communication networks, accurately forecasting the weather, controlling extremely complex machines and factories, safely giving trainee pilots experience of flying an airbus without leaving the ground, etc.

However for this to count as research it must also **increase knowledge**. If it merely uses existing computing knowledge to produce new tools that are useful to increase knowledge about some *other* domain (e.g. physics, biology, etc.) that may make it research in the other discipline. If it increases our **explicit** re-usable knowledge about how to specify, design, build, test, maintain, improve, or evaluate information-processing systems then it is research in the field of software or computer engineering, or AI engineering. (This is not intended to be a precise definition: there may not be one.)

Scientific and engineering research work in Category 3 can be contrasted with a great deal of system development activity that may be of practical use, but either (i) directly deploys existing knowledge in standard ways without extending that knowledge, or (ii) depends only on the intuitive, often unarticulated, grasp of what does and does not work.

As regards (ii), unarticulated intuitive knowledge and skills gained through practical experience (perhaps combined with natural gifts), may be called *craft* since it does not require the use and development of explicit theories about what does and does not work and why (the result of research of Types 1 and 2). Even when such craft work extends what we can do, it is not in itself research and should not be treated or evaluated as such, though it may be a precursor to important research. It may produce useful results but does not, in the process advance communicable knowledge. However craft in building computing systems, like many other types of craft, can, and

⁸Rachel Harrison drew my attention to this point.

often does, later stimulate more explicit science and engineering: we often first discover that we can do something, then later wonder how and seek explanations.⁹ The resulting articulation leads us to understand precisely what was achieved, the conditions under which it can be achieved, how it can be controlled, varied, extended, etc.

3.4 Creation and evaluation of tools, formalisms and techniques to support these activities

Category 4 can be seen as a subset of Category 3, though it may be useful to separate it out because its engineering goals are concerned with the processes of performing the tasks in the previous categories (and this category) and to that extent involves the pursuit of goals which could not have existed but for the existence of computing science. (That's only an approximate truth!)

This category involves a diverse range of activities, including designing new programming languages, new formalisms for expressing requirements, compilers, tools for validating or checking programs or other specifications, tools for designing new computing hardware or checking hardware designs, automatic program synthesisers, tools to support exploratory design of software (e.g. most AI development environments) and many more.

Research on design, analysis and testing methodologies, as well as tools to support them, can be included in this category, though it overlaps with other categories.¹⁰

The design and production of new general purpose computers, compilers, operating systems, high level languages, graphical and other interaction devices and many more, clearly falls into both the third and fourth categories. Moreover, many tools which are initially of Type 4 can migrate into tools of Type 3, e.g. early AI software development tools which were later expanded into expert system shells.

However, it is possible for a tool of Type 4 to have no obvious use outside computing science and yet be of great value. Perhaps an example might be a tool for automatic analysis and checking of the type-structure of a complex formula in a language used only by theorists, or a tool for analysing the structures of complex ontologies developed entirely for research purposes.

3.5 Research on social and economic issues

Research in this category normally requires collaboration with researchers from other disciplines such as psychology, sociology, anthropology, economics, law, management science, political science and philosophy.

It includes attempting to understand all the various ways in which developments in computing technology and artificial intelligence have influenced social, educational, economic, legal and political processes and structures, and ways in which they may influence such processes in the future.

It can also include exposing and analysing ethical implications, including the implications of the impact of the new technology on opportunities, resources, jobs, power structures, etc. for various

⁹It may be that the only way to produce excellent engineers is to start by making them expert craftsmen and women!

¹⁰Tom Addis, Portsmouth University, pointed out in response to the first draft that I had not said anything about research on design and development methodologies. I have now placed this in Category 4, though some aspect of this work clearly belongs in other categories, e.g. exploration of possible methodologies and modelling of human designers. Rachel Harrison pointed out that besides design methodologies there are also analysis and testing methodologies. I have grouped research on all of these together as supporting research of the other types. However, this can also be seen as an aspect of Category 3.

social groups within countries and also the impact on international relations and relative power of nations, international companies, etc. It can also include analysis of ethical implications of views of the human mind arising out of developments in AI.

4 Evaluation Criteria

The five types of research have different evaluation criteria, though there is partial overlap. It is possible that the differences are not fully understood, either by politicians and civil servants who are concerned with funding decisions, or by some of the referees who comment on grant proposals.

In particular where the research is concerned with testing or developing explanatory or predictive theories, the history of science shows that there can be rival theories which are both partially successful and both better than other theories attempting to explain the same phenomena, without there being any decisive way of telling which theory is better, at any particular time. However, As Imre Lakatos (reference above) showed in some cases it is possible over a period of time to tell that one theory is associated with a “progressive” research programme while the other is associated with a “degenerating” research programme.¹¹

4.1 The study of what is possible – and its scope and limits

The criteria for evaluation of this kind of research are subtle, unobvious, and closely related to criteria for evaluation of research in mathematics, logic, philosophy, theoretical physics, theoretical biology, etc. They involve notions like “depth”, “power”, “generality”, “elegance”, “difficulty”, “potential applicability”, “relevance to other problems”, “synthesis”, “integration”, “opening up new research fields”, etc.

It can be very hard for some people who have not done this kind of research to appreciate its value. But there are plenty of widely referenced examples, e.g. Turing’s invention of the notion of a Turing machine and his and Goedel’s work on limit theorems and (less widely known) McCarthy’s invention of a programming language that can operate on expressions in the language – Lisp. (Alas, many developers of programming languages since then have ignored this idea!)

It often turns out that new theories about what is possible also have enormous practical applications, though sometimes these are not understood, or deployable, until many years later. Many deep theoretical advances have had unexpected practical applications after considerable delay. E.g. the problem to be solved may not turn up for a long time, or the application may require additional developments which take a long time: most of the practical deployment of ideas about forms of information-processing had to wait for advances in physics, materials science, manufacturing technology, etc. to produce computers with the power, weight, size, price and diversity of uses that we know today.

Because research of Type 1 is so hard to evaluate and of such potential importance, it may be necessary to devise mechanisms to keep it going and to keep diversifying it with minimal concern for evaluation by *generally agreed* criteria. (Compare the use of stochastic search mechanisms to solve really hard problems!)

¹¹Closely related ideas about the aims of science and evaluation of scientific theories can be found in Chapter 2 of A. Sloman, (1978), *The Computer Revolution in Philosophy, Philosophy, Science and Models of Mind*, Harvester Press (and Humanities Press), Online here: <http://www.cs.bham.ac.uk/research/cogaff/crp> Chapter 2 is: <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/chap2.html>

4.2 The study of existing information-processing systems

Here the criteria for evaluation are more like those in empirical sciences, like experimental physics, biology, psychology, etc. The theories have to be tested against the facts. This can sometimes be done by using the theories to make predictions about behaviour of naturally occurring systems, or by showing how large numbers of different previously observed phenomena can be uniformly explained.

Sometimes theories about natural information-processing mechanisms can be confirmed or dis-confirmed by evidence gained by opening up the physical system or by sophisticated non-invasive techniques for observing internal processes (e.g. fmri scanners).

Often however empirical testing is extremely difficult and has to be indirect, especially when the theory relates to a very complex virtual machine whose structure does not relate in any simple way to the underlying physical machinery, or where the complexity of the physical or physiological mechanisms makes de-compiling an intractable task.

In that case theories may inevitably remain highly conjectural, making it hard to choose between rival alternatives with similar behaviour consequences. Sometimes this leads to sterile debates that would be better postponed until there is a better basis for choosing, while work in the rival camps continues to be supported. Often rival theories cannot be properly compared until long after they are first proposed.

Sometimes, choosing between alternative theories requires introducing very indirect evidence: e.g. showing that the mechanisms of evolution could have produced one sort of architecture but not another, in order to rule out the second as a correct theory of how a human mind works.

But truth is not enough for an explanatory theory to be valuable, for there are trivial or shallow truths: again notions like “depth”, “generality”, “explanatory power”, “elegance”, and a theory’s ability to open up new research problems, are relevant to the evaluation of the theory as a contribution to science. All this is just a special case of philosophy of science, though most philosophers of science are unaware of the special complexities of scientific theories about information-processing systems, because they were brought up to philosophise about simpler sciences such as physics!

4.3 Research involving creation of new useful information-processing systems.

This sort of work has two kinds of criteria for evaluation: how well it extends knowledge and how useful the results are.

Often the work involves both producing new developments of Category 1 or 2 and also deploying them in creating something useful, e.g. exploring ideas about forms of computation, and then later building usable physical implementations of those ideas, or finding a deep explanation of certain diseases then using that explanation in the search for a cure. The two kinds of work need not proceed in that order: in some cases the practical results and explanatory theories may be developed in parallel, or practical difficulties in applying old ideas may point to the need to improve existing theories, formalisms, conceptual frameworks, etc.

In all these cases the criteria for work of Category 1 or 2 are relevant to evaluating work in Category 3 because the work is composite in nature. But there is also evaluation of usefulness of new systems. However, usefulness has its own rewards (e.g. financial rewards) and unless there is also some advance in knowledge it is not research. This must be remembered in evaluating such projects in a research context.

Not everyone will agree on criteria to be used in evaluating practical applications. Most people would agree that results can be evaluated in terms of benefits they bring in enhancing quality of life including new forms of entertainment, or facilitating other activities with important practical goals, e.g. preventing air traffic collisions, allowing secure transmission of confidential messages, or automatically diagnosing skin cancer at a very early stage, or designing a better tool for teaching mathematics. But some people will regard work that builds more powerful weapons that can bring death and destruction (euphemistically named “defence”) as valuable whereas others will condemn such applications. Recent debates about genetically modified food illustrate this point. Moreover, as any *Which?* report shows evaluation can often be multi-dimensional with at best a partial ordering of the options available.

In addition to the evaluation of the costs and benefits of new applicable systems, they can also sometimes be evaluated *intrinsically*, e.g. in terms of how elegant they are, how difficult they were to achieve, how ingenious or original their creators had to be.

Some railway steam engines were beautiful as well as being powerful and fast, and some very useful bridges are also works of art. Lisp (the original version) and Prolog both have a type of beautiful simplicity in relation to their power as programming formalisms, unlike several others I dare not name. Those who attempt to convey a sense of *style* when teaching programming appreciate this point, apart from the fact that a good style can also have practical consequences, such as maintainability and re-usability.

4.4 The creation of tools, formalisms and techniques

These things can be evaluated both according to how well they facilitate work in the other three categories, and also according to the previously mentioned criteria which are independent of usefulness.

Of course producing good tools for doing other things (e.g. for designing and testing models, for building applications, etc.) can be thought of simply as part of those other activities, and evaluated in relation to their indirect benefits. But the good ones have a kind of generality and power that is of value independently of the particular uses to which they are put.

It could be argued that this fourth category is spurious: it should be lumped in as part of the third category, sharing its evaluation criteria. At first I was tempted to do this. However, the development of computing both as science and as engineering has depended on a remarkable amount of bootstrapping, where the most important applications of many tools, concepts, formalisms and techniques are the processes of producing more of the same.

A spectacular example is the role of previous generations of hardware and software in producing each new generation of smaller, faster, cheaper, more powerful, computers.

4.5 Research on social and economic issues

The evaluation of research in this area is a huge topic beyond the scope of this note. However it links up with criteria for evaluating research in all the other disciplines involved in this research, including psychology, sociology, anthropology, economics, law, management science, political science and philosophy. In some cases there are significant disputes about how to evaluate research in these fields and the relevance of those disputes is likely to be inherited by research in this category.