The Question of A.I. Croom Helm, London 1987
pp. 106 - 164

3

Development and Establishment in Artificial Intelligence

James Fleck

3.1 INTRODUCTION

In this chapter, I discuss the role played by scientific establishments in the development of a particular scientific specialty,[1] Artificial Intelligence (AI), a computer-related area which takes as its broad aim the construction of computer programs that model aspects of intelligent behaviour. As with any discussion of a scientific specialty, the identification of what is involved is not unproblematic, and the above serves as an indication rather than a definition. While the term 'Artificial Intelligence' is used in a variety of ways,[2] there is a discernible group (perhaps approaching the degree of commonality to be called a community) of researchers who recognise the term as descriptive of a certain sort of work, and who, if they themselves are not willing to be directly labelled by the term, can locate themselves with respect to it.

Unfortunately, there is little or no commonly available literature that systematically charts the scope of this area. It is worthwhile, therefore, to consider the distinctive socio-cognitive characteristics of research in AI as a prelude to a fairly specific discussion of the social and institutional processes involved in the development of the area,[3] thus providing a basis for exploring the usefulness and applicability of the concept of establishment.

3.2 SOCIO-COGNITIVE CHARACTERISTICS OF ARTIFICIAL INTELLIGENCE

The patterns of research in AI exhibit distinctive characteristics, forming a paradigmatic structure which includes such elements in the scientific activity as research tools, practices, techniques, methods, models, and theories, as well as the normative and evaluative aspects for selecting among them.[4] They serve as guidelines and a basis for future research, but are complexly inter-related, often encompassing contradictory facets in tension.

The elements in the paradigmatic structure of AI are as follows:

(1) The general-purpose digital computer provides an instrumental base and a disciplinary context — computer science — for research in the area. Adequate computing facilities are essential for AI work, and hardware limitations have had a constraining effect. Consequently, the availability of funding has been of crucial importance for the development of the area.

(2) List processing languages, a subset of the high-level programming languages available for exploiting the power of the computer, have been developed as tools for research in AI. The community of people using list processing languages, such as LISP in the United States, or POP-2 in Britain, and their variants, can serve as a first approximation for the AI community.

(3) These list processing languages are orientated towards non-numerical uses, and, hence, contrast with more conventional programming languages such as FORTRAN or ALGOL which are numerically oriented. This non-numerical emphasis, with a focus on logic and structure rather than number, distinguishes AI from areas such as pattern recognition, for example, which depends heavily on the use of statistics.

(4) Associated with list processing languages, there has developed a distinctive body of craft knowledge. A high level of skill, gained through first-hand use and practice, is required for the effective use of any programming language and there are many 'tricks of the trade' as well as distinct programming styles, which can only be absorbed through an extended period of apprenticeship.[5]

(5) Embedded in this craft knowledge are numerous elements such as techniques for problem solving, for representing knowledge, for achieving learning ability,

etc. Particularly important and well developed among these are procedures for carrying out searches, often employing rules of various kinds —heuristics — to guide the search and cut down the possibilities to be explored.

(6) The craft knowledge of AI is deployed in the construction of computer program models — computational models — of some aspect of intelligent activity. These models are generally pitched at the symbolic level of meanings rather than at the physiological level of the underlying mechanisms. This distinguishes AI from many other cybernetic approaches, and from much computer simulation work. The focus on intelligent behaviour provides a disciplinary context — psychology — but, due to the great variety of social interpretations and applications of the term 'intelligence', specific goals for research are not thereby dictated. This lends AI a similarity with what can be termed instrument or technique-based specialties, such as X-ray crystallography,[6] which are free to be applied to various goals.

(7) Associated with the wide variety of specific examples of intelligent activity that have been modelled, a clear research area differentiation has emerged since the early 1960s in which subareas have developed their own particular specialist guidelines and techniques, focused on their own more circumscribed concerns. The research areas that could be identified in the early 1960s were game playing, theorem proving, cognitive modelling (an emphasis on models with psychological verisimilitude), natural language, machine vision, and a range of specific applications,[7] some of which have themselves subsequently differentiated out into well defined research areas. These research areas (or strands of research[8]) constitute a primary setting for scientific activity, and consequently have been one of the basic arenas for competition among practitioners, as will become evident.

These cognitive characteristics, or elements of the AI paradigm structure are, of course, at a very general level. They open up a huge cognitive space and offers wide opportunities for exploration, which were elaborated at a fairly early stage in essentially their complete form, while subsequent work has largely exploited the possibilities opened up. This overview of

the development of AI invites comparison with Edge and Mulkay's account of the development of radio astronomy: the initial discovery of radio waves from space opened up the possibility of a new source of astronomical information — a new cognitive space — which was subsequently exploited by ever more sophisticated methods of detection, leading ultimately to a revolution in the conception of astronomy.[9] However, while radio astronomy was apparently allowed to develop without much external conflict,[10] the same cannot be said of AI.

3.3 COMPETITION AND ESTABLISHMENT IN ARTIFICIAL INTELLIGENCE

Conflict in AI has been bound up with the focus on intelligence. Intelligence is not a socially or cognitively well-defined goal and every distinctive social group tends to have its own implicit definition, couched in terms of its own interests. Consequently research in AI has been oriented towards a variety of goals. This multigoal character leads to a range of struggles between various groups and establishments within and around AI, and is institutionally manifested in the high degree of research-area differentiation, with interdisciplinary and multidisciplinary affiliations, and associated multiple funding sources. This leads to competition on the one hand, between research areas for resources, and on the other hand over the definition of what AI is. This has had quite clear effects on development, as, for example, with the debate in the UK Science Research Council (SRC) in the early 1970s, which led to separate funding mechanisms being set up for cognitive science (linguistics, philosophy, and psychology) applications of AI, and for research within computer science.

This multigoal characteristic, involving competing groups with different aims, has exerted centrifugal pressures on research in the area and has resulted in the non-emergence of a specialty-wide general theoretical dynamic. Attempts at the elaboration of theories of intelligence have informed work in the area — for example, the early programme of research (evident in work in the 1960s in systems such as GPS (General Problem Solver)), towards forming general mechanisms of inference that would embody the essence of intelligence — but these attempts foundered upon the diversity of concepts and applications involved. Later developments such as the attack upon the problem of knowledge representation (a major theme of research in the 1970s) effectively accepted the contingent diversity of intelligence and turned it into a virtue. What theoretical developments there have been, however, have been very specific and localised, often pertaining to the status of the methods and languages employed.

Nevertheless, the absence of a uniform goal or a general theoretical dynamic raises the question of the source of cohesion and co-ordination for research in the area. The answer seems to lie in the craft nature of the paradigmatic structure. While there are many divisions over short- and long-term goals, and between different research areas, there *is* a shared body of technique and practice based on the use of list processing programming languages and transmitted by apprenticeship and personnel migration, thus constraining the historical development of the area. Access to this body of knowledge and skill is restricted by the need for first-hand contact and for adequate computing facilities, consequently leading to tight intercentre and intergenerational linkages in the area. The group of people who control access to these resources clearly constitutes the establishment in AI, and it is at this level that much of the internal research-area competition takes place. It would seem, therefore, that this case demonstrates that scientific establishments need not be characterised by a high degree of solidarity. It would also seem to be the case that a common basis in technique is adequate to hold an area together in the face of strong centrifugal tendencies, especially where it is associated with an instrumental basis for research in the area. The need for adequate computing facilities has restricted access to the field and encouraged the development of a strong communication infrastructure, particularly in the United States where the ARPA (Advanced Research Projects Agency) computer network enables researchers at geographically distant sites to communicate as easily as if they were at the same location. Such expensive instrumental needs have, of course, opened the area up to influence from funding agencies. Research in AI has without doubt depended upon substantial support from various agencies, such as the US Department of Defense, and the UK Science Research Council, and has consequently been shaped by the concentration policies of these agencies which have had the effect of consolidating the position of the establishment in

the area. However, it is difficult to find evidence for any positive direction of research by the funding agencies during the 1960s, although in the stringent financial climate of the 1970s this did change, and tighter demands for the attainment of particular goals were made, resulting in a restructuring of funding patterns among AI centres, both in America and in Britain. Moreover, during the early 1980s, there appears an initial emergence of 'AI-technology', where specific lines of research, considered to have commercial potential (not necessarily those of prime scientific importance), are being picked up out of the university context, along with supporting personnel, and transferred for industrial development. At this level of a broad overview of AI, there appear to be some similarities to Yoxen's description of molecular biology in terms such as 'directed autonomy'.[11] However, at the more detailed level of the following discussion, it is very hard indeed to identify elements of long-term direction that might fit in with concomitant long-term strategies on the part of the funding agencies. Perhaps what is at issue here is the appropriate size and nature of the envelope within which autonomy is exercised by the practitioners, while yet remaining suitably circumscribed in accordance with the externally imposed direction.

The availability of funding and institutional resources for the area as a whole is, of course, controlled by a wider establishment — the funding agencies and the universities — and in the following discussion the processes of negotiation between the specialty and wider establishments stand out clearly. It will become clear that the response of the wider establishment of AI is by no means uniform, thus illustrating that too monolithic a character should not be imputed to the establishment at this level either, but that an adequate understanding needs to take into account the particular features of the area — that is, the specificity of AI.

The allocation of resources and the negotiation between the establishments in and around AI have been clearly affected by what Elias termed a struggle for monopolisation of the means of orientation.[12] Research in the area is often seen as constituting a further thrust of mechanical materialist science into an area — the nature of mind — hitherto under the exclusive sway of traditional cultural values. By and large, mind or intelligence is regarded as the most characteristic and unique of human attributes — extremely rich and complex, and undoubtedly beyond the reach of scientific analyses more appropriate for the understanding of inert matter. Thus, the focus on intelligence and mind brings AI into an arena of conflict at a deep-seated emotional level which touches immediately upon everyone's image of themselves, and induces strong 'for and against' alignments. It is an issue of general public rather than narrowly scientific interest, as is evidenced by the high relative exposure AI receives in the press, on television, etc. It is doubtful whether many other scientific areas could have the same effect — with the exception of some areas such as genetic manipulation which undoubtedly bear comparison. The nature of mind is an area where the religious and philosophical establishments still claim authority, and AI has to fight for legitimacy. Even where explicitly religious commitments do not seem to be involved, those brought up under the Western humanist culture often feel threatened by what they see as the reductionist nature of AI, and work in the area has been denounced as bad science, non-science, gross reductionism, and even immoral science. At this level, then, there is negotiation and conflict with establishments outside science, as well as between establishments within science.

As far as competition between establishments within science is concerned, the case of AI illustrates an important, characteristically twentieth-century development in scientific thinking — the software sciences — with a focus on pattern and organisation rather than on the properties of substance or matter. Yoxen points out in his paper the importance of such metaphors as code, information, read-out, program, etc., in the reconstitution of biology: with AI such ideas are at the very core of the subject. Moreover, with the increasing penetration of the computer into all areas of science and scholarship, the features of AI related to the software science nature of research in the area, may well become typical of many fields of science. In particular, the diffuse, method-based character of AI, with its contingent adaptation to diverse substantive issues, poses a contrast and challenge to the coherent, theoretically centred nature of the current scientific ideal, deriving from the example of the dominant physical sciences tradition. The former would not seem to support so readily a monolithic unified establishment as does the latter, and the former may consequently have implications for the future development of the sciences.

Thus, AI appears to be an interesting case in the context of a discussion of scientific establishments, for a number of reasons. The more diversified nature of establishments in the software sciences may have wider implications for the sciences as a whole; negotiation and conflict between establishments at a variety of levels is clearly illustrated — including a struggle for the monopolisation of the means of orientation; the power bases of these various establishments (control over cognitive, instrumental or financial resources) are clearly evident, and finally AI provides many examples of the problems arising from struggles between the various establishments involved in a multi-disciplinary and interdisciplinary area of research — illustrating many of the points commented upon by Elias.

3.4 EARLY DEVELOPMENT IN THE UNITED STATES

The Second World War acted as a melting pot for various quite different lines of research and disciplines. In the intense concentration on the common goal of winning the war, traditional disciplinary boundaries were broached and new areas of research emerged, such as information theory, operations research, cybernetics and, of course, the development of the digital computer itself. These areas of research can be broadly characterised as the software sciences, in that they focused on pattern and organisation rather than on substance or matter — the concern of the natural sciences such as physics and chemistry.

Cybernetics, a rather general field given a name and identity by Norbert Wiener's classic book, *Cybernetics — control and communication in the animal and the machine*, was concerned with the essential similarities between machines and biological processes.[13] Work in the area developed during the 1940s, and involved such approaches as the comparison of biological and neurophysiological processes with electrical circuits and networks of artificial neurons, or the investigation of the general principles of adaptation in self-organising systems — systems which were rich in feedback connections, and would settle into stable configurations after being disturbed.[14]

The advent of the digital computer in the early 1950s heralded a new approach which sought to build models of intelligent processes at the symbolic level.[15] Concepts were represented and operated upon directly in the computer using high-level programming languages. These 'symbolic' models represented intelligent activity at the level of thought itself, rather than at the level of the physiological mechanisms underlying thought, thus contrasting sharply with other cybernetic approaches.

In 1952, a conference was held under the rubric 'Automata Studies'.[16] This conference, organised largely by John McCarthy, was intended by him to attract proponents of the symbolic modelling approach. It failed in this aim, and attracted contributions more clearly in the other cybernetic traditions. This determined McCarthy to 'nail the flag to the mast the next time', which he did by explicitly using the term 'artificial intelligence' in a subsequent summer school held at Dartmouth College in the United States in 1956, to discuss 'the possibility of constructing genuinely intelligent machines'.[17] The official title was 'The Dartmouth Summer Research Project on Artificial Intelligence' and it did succeed in isolating the symbolic modelling theme. Among those present were J. McCarthy, M.L. Minsky, H.A. Simon and A. Newell.[18] After the meeting, Simon and Newell were to start a group at the Carnegie Institute of Technology (now the Carnegie-Mellon University) with the aim of developing models of human behaviour, while McCarthy and Minsky built up a group at the Massachusetts Institute of Technology (MIT), with the goal of making machines intelligent without particular reference to human behaviour. Later in 1962, McCarthy was to move to Stanford University, where he initiated another AI project. These three centres, along with Stanford Research Institute, dominated AI research in the United States in the 1960s and 1970s. Also present were C.E. Shannon (known in AI for his outline of the chess-playing paradigm which is essentially the same as the one underlying the microelectronic machines that can now be bought off the shelf) and A.L. Samuel (who developed an impressive checkers playing program which incorporated an elementary learning mechanism).

It was at that meeting that the broad outlines of a distinctively AI approach — indeed, what might be called a proto-paradigmatic structure — emerged. This involved the use of high-level programming languages to provide symbolic models of various aspects of intelligent activity. The first areas attacked, chosen partly because they seemed to epitomise intelligence, and partly because they were sufficiently well

defined to be readily programmable, were theorem proving in mathematical logic, and games such as chess and checkers. While chess and other games employed numerically based techniques for choosing board moves, the 'Logic Theorist' of Simon and Newell (which was presented at the Dartmouth conference, the first working, characteristically AI program developed) employed non-numerical techniques.[19] During the late 1950s, programming languages designed specifically for non-numerical symbolic information processing were developed by those two researchers, along with J.C. Shaw,[20] and in 1960, McCarthy formulated LISP (list processing language) which became, and still is, the most widely used AI language.[21] During the late 1950s, it also became clear that organisational techniques of search were of paramount importance in attaining the desired ends, and that the numerical aspects were of secondary importance. The principle of looking for and using certain heuristics — that is, rules of thumb which might help in finding a solution but which would not guarantee a solution — became established.[22] By the early 1960s, various successful programs had been written, resulting in a general air of optimism, and indeed by this time the paradigmatic structure of AI had been elaborated in essentially its complete form, as already described.

3.5 THE ESTABLISHMENT IN THE UNITED STATES

At first sight it might seem remarkable that McCarthy, Minsky, Simon and Newell, without doubt the four 'great men' of the AI establishment, should all have been present at the Dartmouth meeting. I would argue, however, that the emergence of the American establishment in AI was part and parcel of the process of defining the paradigmatic structure of research in the area and the organisational structure of the field.

In the first place the four were actively involved in the organisation of the field. McCarthy, as already noted, arranged meetings to bring together those interested in the very loosely defined goal of constructing genuinely intelligent machines. Subsequently, he went on to found a group at MIT, along with Minsky (a fellow student with him at Princeton), and later the group at Stanford. Simon and Newell developed the group at Carnegie.

In the second place, these four men were centrally involved in defining the substantive cognitive elements of the AI paradigmatic structure already outlined. Simon and Newell, as well as presenting the first working AI program, the Logic Theorist, had also developed a series of list processing languages, the IPL series, the forerunners of the basic element of the AI research activity. McCarthy had produced the definitive AI programming language, LISP, on the basis of these forerunners, and had also incorporated certain features which embodied the emphasis on logic rather than numerical mathematics. Minsky had written an influential systematising paper which explicitly outlined the importance of heuristic search. As well as producing the basic tools for subsequent research, these four also defined in broad terms higher-level guidelines for future research. Simon and Newell pioneered the focus on investigating human cognitive processes as a source of inspiration for computational models, while McCarthy and Minsky went in more for the idea of investigating mechanisms for achieving intelligent activity in the abstract, without prejudice towards specifically human forms.

The fact that the four were involved in the formation of AI research activity from the beginning has contributed to their success in becoming members of the establishment in two ways. In the first place, the founders of any field, simply because at that stage they are competing with fewer people, gain a visibility which later contributors are unlikely to attain, unless they in turn can produce work that will lead to subsequent distinctive and fruitful development. In the second place, all four have been around for a long time, and consequently the process of normal scientific career progression has ensured their continued visibility and position within the establishment.[23] Moreover, once recognised as members of the establishment, they have in fact continued to be influential within the field. McCarthy's suggestions for a mathematical theory of computation,[24] and his emphasis on the use of the predicate calculus, have been themes which have been taken up and developed. Minsky has continued to produce influential synthesising research programmes, as for example in his presentation of the theme of semantic information processing,[25] or more recently with his explication of 'Frames' — high level data structures for organising and mobilising the vast knowledge bases with which effective AI programs have had to work.[26] He has also used his establishment position to argue effectively against the formal theorem proving strand of research.[27] Simon and Newell have continued to pioneer new approaches at the psychological interface of AI and the 'production systems' formulation developed out of their earlier work on general problem-solving, and promoted by Newell, has been widely taken up.[28]

However, the members of the AI establishment did not arrive from nowhere. Their success, without doubt, owed much to their having attended prestigious institutions as students, and to their being sponsored by people who were already members of the wider scientific establishment, not necessarily in the cybernetic area. McCarthy, Minsky, and Newell all attended Princeton as graduate students, for instance, and McCarthy worked for Shannon on the organisation of the 1952 Automata Studies conference; while Minsky was associated with W. McCulloch, whose 1943 paper with Pitts is recognised as another of the texts marking the emergence of cybernetics.[29] Simon had already established his reputation in the fields of political science and economics,[30] and he himself acted as sponsor for Newell.

Moreover, this intergenerational establishment reproduction process continued, and students of McCarthy, Minsky, Simon, and Newell have largely dominated the field. Figure 3.1 gives a graphic illustration of the prevalence of such links — links which have led to charges of nepotism being levelled at the AI establishment.[31] This structure of very strong intergenerational linkages turns out to be characteristic of the development in Britain as well, and in the section on the establishment in the United Kingdom, some underlying reasons for the strong linkages are discussed.

The emergence of this group as the establishment in AI was undoubtedly consolidated by its success in getting the backing of the US Department of Defense, mainly through the Advanced Research Projects Agency (ARPA), which provided some 75% of US AI funding for the ten years from 1964, and through the Air Force.[32] Furthermore, the preference on the part of the ARPA for concentrating resources in a few selected centres guaranteed the position of the establishment, especially in view of the great expense of adequate computing facilities, which effectively barred other groups from competing.

Another aspect of development in AI which has been characteristic and of importance for the field, and which has served further to reinforce the position of the establishment,

Figure 3.1: The establishment in the United States, 1960—mid-1970s. The members of the establishment were derived from a consideration of the editorial board of the *Artificial Intelligence Journal*, conference organising committees, invited conference speakers and panel members, supplemented by well known researchers as judged on the basis of a reading of the literature. They include 73 out of a total of upwards of 500 contributors to the area. Available data were limited, but indicated that only some 11 out of the 73 had *not* worked at some time or done a PhD at one of the big four AI centres: Massachussetts Institute of Technology (MIT); Carnegie Mellon University (CMU); Stanford University (SU); and Stanford Research Institute (SRI). At least 24 of the 73 received their doctorates from one of MIT, CMU, or SU. There is no particular significance in the ordering, nor is the record of movements complete. Intergenerational and intercentre linkages are probably underestimated due to lack of data.

( ) ⟶ movement of personnel.

⟶ student of/worked for.

⟷ colleague relation

═ boundary between the wider establishment, and the Artificial Intelligence establishment.

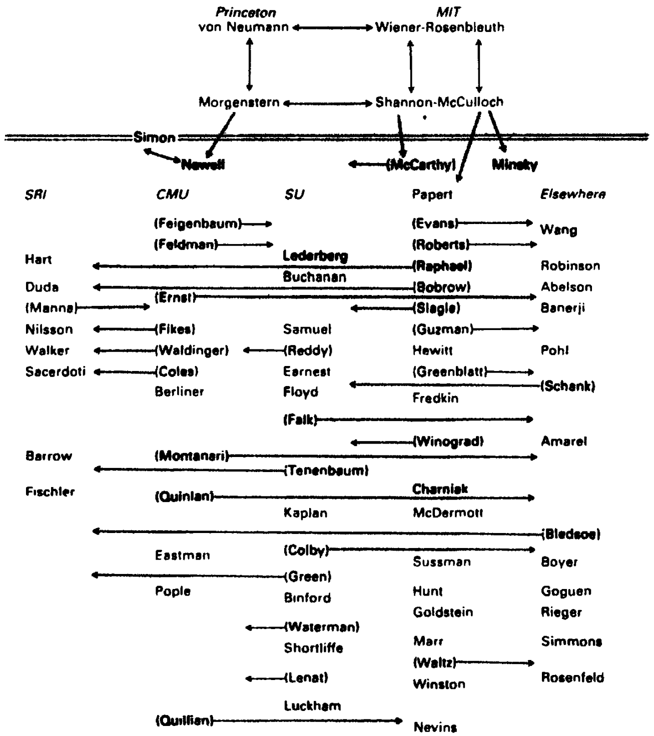

should be noted: this is that the general aim of research in the area to produce intelligent machines has excited extreme reactions and has tended to lead to very strong for-and-against[33] alignments. Such a reaction is not at all surprising given the sensitivity of such a goal to peoples' images of themselves. Here Elias's comments about the competition for the monopoly over the means of orientation are relevant.[34] The AI approach is seeking to establish and legitimate a view of intelligence and the nature of mind which challenges the received commonsense view of mind and intelligence as something rather special and certainly well beyond the reach of scientific analysis. Moreover, this received view is very much under the sway of the religious establishments, or where religious authority does not hold, under the sway of a liberal humanist tradition. Strong reactions are commonplace in AI and, on the sociological level, have probably had the effect of heightening the difference between those on the inside and those on the outside, consequently reinforcing and concentrating the position of the establishment.

Thus, it can be observed that the emergence of the American establishment was very much bound up with the development of AI as a distinctive area of research, and its position was consolidated by the success in gaining backing from the Department of Defense. The American establishment was not only involved in providing an organisational basis for research in the area, but was also very closely concerned with the elaboration of a distinctive cognitive basis for research in the area, the AI paradigmatic structure. In the following discussion of the development of AI in the United Kingdom, some of the themes already introduced will be reiterated, while other issues will become evident.

3.6 DEVELOPMENT IN THE UNITED KINGDOM

In Britain during the 1940s, there was a similar flourishing of interest in general cybernetic concerns as occurred in the United States, and discussions of the possibility of machine thought were common.[35] A.M. Turing was an enthusiast for the possibility of intelligent machines, and his 1947 and 1950 papers still stand in many respects as definitive surveys of the arguments for and against AI.[36] R.J.W. Craik, whose 1943 book *The nature of explanation* is recognised as one of the texts marking the emergence of cybernetics, wrote passages that bore a remarkable foreshadowing of the actual AI paradigmatic structure, as for example in the following passage:

> . . .thought models, or parallels, reality — that its essential feature is not 'the mind', 'the self', 'sense data', nor propositions but symbolism and that this symbolism is largely of the same kind as that which is familiar to us in mechanical devices which aid thought and calculation.[37]

Others were interested in more specific cybernetic approaches: W.R. Ashby, whose name is perhaps second only to N. Wiener's in association with cybernetics, contributed many ideas and books on the subject;[38] W.G. Walter, who achieved a degree of fame with his electronic tortoises which exhibited elementary reflexive behaviour;[39] F. George, who wrote on a cybernetic approach to the brain;[40] and A.M. Uttley, D. Mackay and others, who worked on neural net models of cognition and perception. Furthermore, through the informal RATIO club which existed during the early 1950s (to which Turing, Mackay, Ashby, Walker, Uttley and others belonged) there was frequent interaction and discussion on these issues — discussion which often also involved researchers from the United States: McCulloch, for example, attended the first meeting of the club in 1949.[41]

These discussions continued through the 1950s. There were several British contributors to the 1952 Shannon and McCarthy Automata Studies meeting, and in 1958 a conference on the 'Mechanisation of Thought Processes' was held at the National Physical Laboratory in England, to which people such as McCarthy and Minsky contributed, as well as the proponents of other cybernetic approaches and practitioners of the art of using the digital computer.[42] Indeed, it was there that McCarthy presented his 'Advice Taker', a suggestion for using the predicate calculus for modelling commonsense reasoning, an approach which proved influential within the more specific AI context.[43]

Nevertheless, despite all the activity and discussion in the area, the specifically AI approach was not developed in Britain until the mid-1960s, and then the prime mover was someone quite external to the cybernetic network: this was Donald Michie, a geneticist at Edinburgh University. During the War he had worked with Turing at Bletchley and became fascinated

in the possibility of constructing machines that could think.[44] There were no opportunities for him to follow up these interests after the War, except on the hobbyist level, so he took medical sciences at Oxford and subsequently specialised in genetics.[45] In 1962, however, during a visit to the United States (arising out of his hobbyist work on trial-and-error learning in the game of noughts and crosses) he became aware of the developments in AI there and was impressed by the computer facilities available. On his return from the United States, depressed by the lack of such facilities in Britain, he became active in agitating for something to be done — lobbying, writing newspaper articles, and so on.[46] At about the same time, he had grown dissatisfied with his position in the Department of Surgical Science at Edinburgh, where he held the post of reader, and in 1963, very much on his own initiative, he moved out of the department and with his secretary and part-time helpers set up an unofficial unit — the Experimental Programming Unit. Meanwhile, he continued his lobbying for improved computer facilities and had buttonholed C. Jolliffe, the deputy grants director of the then Department of Scientific and Industrial Research (DSIR), whom he impressed with his concern over the state of UK computer provision, and was consequently commissioned to carry out a survey of computing interests and views among British scientists.[47] Lord Halsbury had just been charged with setting up a specialised computer board in the DSIR (subsequently the SRC) with the aim of reviewing and supporting research in computer science. The results of Michie's report indicated a widespread positive assessment of the potential of AI among young computer scientists and this undoubtedly formed the basis for the proportionately generous funding by the computing board of the SRC for research in the area during the mid- to late 1960s.[48] This took place against the background provided by the Flowers Report of 1966 on computing in universities and colleges, which recommended a large expansion in computer provision and training, and set the outlines of the presently existing regional computing centre structure, based in London, Edinburgh, and Manchester.[49] In this context, Michie, with his energy and enthusiasm, was seen very much as a bright young man, and he succeeded in attracting several large grants,[50] which no doubt encouraged the university authorities to give official recognition in January 1965 to his irregularly established unit, which grew rapidly over the next few years.

It is interesting at this point to consider the substantive lines of research that emerged at Edinburgh,[51] for the research profile bore a remarkable similarity to the patterns evident in the United States. Michie's own immediate concerns were with game-playing and heuristic search, then regarded as central to the field. Also, in the Experimental Programming Unit there was a project involved with human problem-solving studies, similar in flavour to the work of Simon and Newell at Carnegie; and there were also projects concerned with developing the instrumental base for AI research. One element was the Mini-Mac project, an interactive multi-access system (and the second such system to be developed in Britain), so called after the first project of its kind, carried out on a grander scale — project MAC at MIT, and also partly motivated by AI concerns, according to McCarthy.[52] Another element in the instrumental base was the development of a list processing language, POP-2, which became the staple AI language in the United Kingdom (and with refinements is still in use today),[53] just as LISP was the staple in the United States.

An independent group worked in the Metamathematics Unit with a focus on automatic theorem proving. This unit had been set up by Bernard Meltzer, a reader in the Department of Electrical Engineering, who had become interested in the use of the computer in the course of his work and combined this interest with his hobby of mathematical logic. Like Michie, he took the fairly dramatic step of moving out of his official department and into a new, unofficial unit after a visit to the United States, where he had visited similar projects.[54] Due to their common use of symbolic rather than numerical information processing, and the intelligent nature of their goal, theorem proving, these approaches had developed as an autonomous but central strand in the AI area of research.

Thus, by the mid-1960s, there had emerged in Edinburgh a centre for research in AI which reproduced the same features as had developed in the United States. In September 1965, Michie organised the first of a series of meetings — the Machine Intelligence Workshops — which were held in Edinburgh and attracted leading AI researchers from the United States, as well as interested people from elsewhere in Britain.[55] These workshops played an important part as a forum for discussion and the communication of the AI approach, and influenced such people as E.W. Elcock and J.M. Foster of Aberdeen University, where in an SRC-sponsored computer unit they worked on game-playing programs and high-level programming languages and systems which incorporated some AI elements.[56] M.B. Clowes was another interested researcher who attended. Encouraged by Michie, he had set up AISB (the Society for Artificial Intelligence and the Simulation of Behaviour) in 1964, which, after an informal and hesitant start has grown into a thriving learned society for the area, holding two yearly conferences and publishing a regular newsletter.[57] Clowes had also visited the United States and had been impressed by what was happening there. He did not work on specifically AI projects in the mid-1960s, and depressed by the lack of suitable computer facilities, he went to Australia for several years. However, he had met and impressed N.S. Sutherland, who had been interested in mechanistic models since the 1950s, and who built up a centre in experimental psychology at the newly established University of Sussex at Brighton in the 1960s. On his return from Australia, Clowes went there to work in AI on a grant held by Sutherland. This work heralded the emergence of Sussex as a major centre for AI in the 1970s.

It was during meetings of the AI community following the technical business of the workshops when suggestions were first put forward for starting a speciality journal for the area — *Artificial Intelligence — An International Journal*, which was eventually founded in 1970 and of which Meltzer became the editor.[58] It also appears that suggestions for establishing international conferences on AI were discussed at Machine Intelligence Workshops:[59]the first was held in 1969, and since then these conferences (held every second year) have grown steadily in size and importance.[60]

Thus, it is clear that the Machine Intelligence Workshops were of great importance for the social development of the field at the international level and consequently, firmly establishing Edinburgh as an AI centre of international reputation. This reputation was further enhanced when Michie succeeded in attracting to Edinburgh from Cambridge the research group of Richard Gregory, the psychologist who became known for his book on the eye and brain,[61] and H.C. Longuet-Higgins, a theoretical chemist of international standing. The basis for the merger was the goal of building an intelligent robot — a goal that was also being pursued at other major research centres in

AI in the United States: MIT, Stanford University, and Stanford Research Institute.

The robot project was seen to pose a challenge for AI in that it required the integration of many of the strands of work within the area — machine vision, problem solving (often based on theorem proving methods), manipulation of a hand in three-dimensional space, and even natural language for communication. Gregory's group was to provide the perception and engineering aspects, while Michie and Longuet-Higgins worked on the problem solving and cognitive aspects. For this project the SRC awarded a major grant and provided a new computer. In addition the Nuffield Foundation provided funding for equipping an engineering laboratory to build the robot and associated hardware and other devices.[62]

Part of the deal involved in attracting these research groups was that the University of Edinburgh, largely as a result of the good offices of Michael Swann (now Lord Swann), then the Vice-Chancellor, would set up a new department — the Department of Machine Intelligence and Perception — and provide chairs for the new senior people involved.[63] This was the first specifically AI-focused department anywhere in the world, and was seen by some as an exciting venture, though others were less welcoming. In particular, as an independent institutional entity in the university, it was in direct competition with the new computer science department, resulting in rather distant and at times antagonistic relations between the two departments.[64] This contrasted with the situation in America, where AI was usually carried on within the departments of computer science and electrical engineering.

This institutional innovation would not have been possible had it not been for two favourable factors. The first was the general context of university expansion of the early and mid-1960s. This expansion, in fact, started being curtailed just after the establishment of the new department, and the resulting squeeze contributed to the problems which beset the department, as the level of University Grants Committee support could not keep pace with the large Research Council funding that the new, rapidly growing area attracted. Had the curtailment of expansion come a few years earlier, it is highly unlikely that a separate department would have been approved. The second factor was the support afforded by Swann: as Dean of Science he had backed Michie's previous initiatives, and newly

elected in January 1966 as Principal of Edinburgh University, he was very receptive towards new departures and constantly promoted the status of Edinburgh as second only to Cambridge in research.[65] Without such sponsorship, it is again doubtful whether a new department would have been instituted, or whether Michie would have succeeded in attracting Gregory and Longuet-Higgins. However, the department was established in October 1966, and while in the event the institutional attractions were evident, it is interesting to consider the scientific motivations for these people with a non-AI background to change their area of research.

Gregory had engineering interests which led him to seek a new methodology involving a closer study of the physical basis in the brain for perception and cognition than was usual in psychology at that time, and the AI approach seemed to promise developments along these lines.[66] He also brought with him other members of his group, notably S.H. Salter, who had extensive engineering competence and built the robot hardware (and was later to become known for his wave-power system — the Salter duck); and J.A.M. Howe, a psychologist who was to explore the applications of AI in educational research, and who became head of the AI department at Edinburgh in the late 1970s.

Longuet-Higgins was very much a scientific high flyer, achieving international distinction in his work in theoretical chemistry with C.A. Coulson at Oxford, and gaining a professorship at the early age of thirty. For his eminence in chemistry, in 1958 he was elected a Fellow of the Royal Society, and in 1968 a Foreign Associate of the US National Academy of Science, the highest American honour available to a non-US citizen. Despite his great success in chemistry, or perhaps because of it (in that he was motivated to seek similar success in a new and potentially exciting but unexplored field), he had joined with Gregory in planning a Brain Research Institute. Negotiations were in hand for funding from the Nuffield Foundation, and for accommodation at Sussex University, when Michie, at a meeting with Gregory early in 1966, suggested that Edinburgh would be an ideal centre in view of its already established AI work. With the institution of the new department at Edinburgh, Gregory's group and Longuet-Higgins moved to Scotland, and great hopes were entertained for the future of co-ordinated research in the area.

The anticipated co-operation failed to materialise. Problems over accommodation, personal, political, administrative, and scientific factors were involved in what became a very complex and confused situation during the late 1960s and early 1970s. Gregory never really settled in at Edinburgh nor became involved with the computational approach, though he remained favourably inclined towards it, and in 1970 he finally left to go to Bristol University. The engineering workshop in the Bionics Laboratory had proceeded, however, with the building of the robot hardware, and a prototype was connected to the computer for the first time in May 1969. Longuet-Higgins did absorb the computational approach, but irreconcilable differences between him and Michie over the installation of the new computer and the robot project, as well as over their approaches to work in the area, soon emerged and resulted in Longuet-Higgins moving into separate accommodation and thenceforth running his unit (then called the Theoretical Section) quite independently apart from access to the common facilities.

Michie favoured a rather swashbuckling style of directing large team projects oriented to goals which could be linked with industrial applications and, in fact, was involved in launching a university-based company to market compiler systems and other software for the POP-2 language, which was developed in the department.[67] In addition, he was extremely energetic and persuasive and very successful in obtaining funding from many different sources.[68] Longuet-Higgins, in contrast, favoured a more restrained, academic style, preferred an individual basis of working with a few colleagues on research chosen purely for its intrinsic scientific interest, and was dubious about the advisability of mixing commerce and industry with research. These differences in style, aggravated by contrasts in personality, were associated with conflicting views on AI: Longuet-Higgins thought that 'artificial intelligence' was not a science or technology in its own right, but was a new way of tackling problems in those existing sciences which were relevant to the phenomena of intelligence. It set new standards of precision and detail in the formulation of models of cognitive processes, those models being open to direct and immediate test.[69] Michie's position, on the one hand, was closer to the view

> . . .that success in achieving the long-term aims of Machine Intelligence should be regarded as the major goal of Computing Science. Furthermore, progress in Machine Intelligence is continually generating pressures for solutions to fundamental problems of Computing Science in an environment where they will be used; an environment which by its very nature, demands quality and generality. As a result a decision to invest apparently disproportionate sums into Machine Intelligence could only have beneficial effects to the whole of Computing Science.[70]

This lack of consensus among the practitioners within AI was undoubtedly a complicating factor when the SRC Computing Science Committee came to review its funding policy in the early 1970s. The 'cognitive science' view of AI put forward by Longuet-Higgins was not seen by the reviewing panel as falling within its scope, while there was strong opposition to the view that machine intelligence should be regarded as a major goal of computing science. These differences in views were noted in the SRC *Computing Science Review* of 1972, and were given by Sir Brian Flowers, then chairman of the SRC, as being among the reasons for commissioning Sir James Lighthill to review the field,[71] a review that was to have a large impact on the area as will be discussed in due course.

The differences between Michie and Longuet-Higgins also caused great internal problems at Edinburgh, resulting in frequent appeals being made to the Principal and Secretary of the university, and to the SRC, and a bewildering sequence of organisational forms were instituted by the university authorities in attempts to alleviate the embarrassing situation; but to no avail. Michie, however, started losing the support of his other colleagues, and finally, in 1974, following a lengthy and, to Michie, unsatisfactory, review of the situation,[72] a new Department of AI was set up with Meltzer at its head and comprising most of the AI research groups in Edinburgh. Michie was given his own independent Machine Intelligence Research Unit. The bulk of the very considerable resources which had been built up over the decade, including the robotic equipment, was settled with the department, and, furthermore, limits were placed upon the scope of future research efforts by Michie.

However, despite the tension between the senior people at Edinburgh, there was a thriving research environment in the late 1960s and early 1970s, with frequent visits by people from elsewhere in the world, including the United States and the up and coming Japanese AI-oriented groups.[73] Young researchers were very successful, sometimes gaining an international reputation before obtaining their doctorates: P.J. Hayes, for example, became well known after publishing a joint paper with J. McCarthy in 1968,[74] only receiving his PhD in 1972. G.D. Plotkin's work on inductive inference was considered outstanding;[75] R. Kowalski made a name for himself with his vigorous promotion of the predicate calculus as a programming language in its own right, an approach which became an independent strand of research termed 'Logic Programming';[76] while R.M. Burstall, one of the original members of the Experimental Programming Unit, and Michie's second in command, became established as an outstanding computer scientist with an international reputation in his specialist area — the theory of computation, in which he built up his own group in the 1970s. In 1978, he was appointed to a chair with the title 'professor of AI', despite the fact that computation theory was by that time a general computer science research area, rather than a specialist AI one. The robot project attracted considerable publicity with some five television and film crews visiting:[77] indeed, demonstrations became so frequent as to interfere with the everyday research work and had to be restricted.[78]

However, the concentration of talent, the surfeit of publicity, and perhaps more than anything else, the predominance of research over teaching in AI, attracted hostility from other departments weighed down with heavy teaching responsibilities,[79] and strong pressures grew for the area to normalise its activities. In addition the lack of a career structure for researchers on short-term contracts, coupled with the increasing uncertainty over the future of the centre due to the leadership tensions, led people to start moving elsewhere: Hayes went to Essex; Kowalski to Imperial College; several other researchers to the United States; and Longuet-Higgins himself, along with members of his group, moved to Sussex University. In the course of a couple of years, therefore, many of the most highly respected researchers left Edinburgh, and in some quarters Edinburgh was viewed as being in decline.[80]

The problems facing AI were not restricted to Edinburgh alone, nor was the division of the department the outcome of purely local politics: rather, these events were tied up with national attitudes especially on the part of the SRC. The SRC

had never been happy with the breakdown of co-operation over the major robotics grant and had become dissatisfied with the progress made in work on the project. This dissatisfaction stemmed to some extent from a basic lack of sympathy with the goals of those involved in the project. Michie's very ambitious plans for a seven-year industrially oriented programme of research in robotics failed to win favour with the SRC and was never formally submitted.[81] More modest proposals were put forward to maintain the level of effort on robotics, but despite a year's very intensive work on the project, in which programmable assembly using visual recognition of parts was attained[82] (at that time one of the foremost achievements in robotics in the world, comparable with leading work in America and Japan), these proposals were turned down. At that stage, the SRC had become very impatient with Michie, as his entrepreneurial talents did not fit in with their expectations, and the previously mentioned survey of AI by Sir James Lighthill, FRS, Lucasian Professor of Applied Mathematics at Cambridge, and an eminent hydrodynamicist, undoubtedly influenced their decision.

Lighthill's report created a major controversy and was published in April 1973, along with other assenting and dissenting views. Lighthill was highly critical of AI in general and suggested that there were three basic categories of research in the area: work aimed at advanced automation on the one hand, and at computer-based central nervous system research on the other, with in addition a bridge category with the basic component of building robots, which he saw as the essential underpinning for AI to have any claims to unity and coherence. Progress in this category, Lighthill suggested, was virtually non-existent and the building of robots a mistaken enterprise, possibly motivated by a desire on the part of those concerned to 'minister to the public's general *penchant* for robots by building the best they can', and possibly also by 'pseudomaternal' drives to compensate for male researchers' inability to give birth.[83] (It is not hard to detect a reference to Michie's polemical enthusiasm in these comments.) Furthermore, what success there had been, he suggested, was evident only in particular applications and derived from knowledge contributed from the substantive fields modelled, rather than from any AI component. In time, he saw the bridge category as withering away, while work directed towards the two extremes would become integrated with other research in their general areas.

Not surprisingly, this caused a major stir in the AI community across the world,[84] and the resulting controversy received much public airing in the press and even on television.[85] Without a doubt, despite Lighthill's protestation that his report

> would simply describe how AI appears to a lay person after two months spent looking through the literature of the subject and discussing it orally and by letter with a variety of workers in the field and in closely related areas of research,[86]

it delivered a blow to the prestige of research in the area from which it has never fully recovered. While Lighthill's comments on robots were directed at the specifically AI category, it appears that they also had some effect on inhibiting robot research and use in Britain in general,[87] whereas in other countries robotics has been a steadily expanding area throughout the 1970s.

In practical terms, the report did affect financial support for research in AI in Britain, particularly in the case of Michie's proposals for robotics research, and it also had some influence on funding in the United States where AI robotic projects were cut back, and the Advanced Research Projects Agency (ARPA), the main sponsor of American work in the area, started insisting on mission-oriented direct research, rather than basic undirected research.[88] These cutbacks took place in the context of the general reduction in public spending in the early to mid-1970s, which affected scientific research in all areas, especially those not seen to be of 'social relevance'. However, the effects were to some extent mitigated in the case of AI: partly by the variety of funding sources supporting the area; partly by the SRC's identification of machine intelligence as an important area of long-range research in its 1972 *Computing Science Review*, which underlined the fact that there was, in any case, no one predominating view on the value of AI; and partly by the expanding nature of computing science in general. Consequently, particular projects were able to get support, especially if their relevance was emphasised and explicit reference to robotics avoided.[89]

Furthermore, the debate over the Lighthill report also led to cognitive science being recognised by the SRC, and a panel was set up to review applications in this area. Thus, there was to some extent a shift in resources to this area, rather than a straightforward cut-back of AI as a whole.

That the reorganisation of AI at Edinburgh, with effective removal of Michie from a central position, was not a purely local affair, but was rather bound up with changes in attitude in the wider scientific establishment, given expression by Lighthill, is borne out by the similar pattern of events occurring at Aberdeen.[90] There, the computer unit was dissolved in 1972 after the SRC refused to renew its grant. Elcock, the organisational prime mover behind the AI interests there (and a colleague of Michie's) had some differences of opinion with the university authorities and the newly established computer science department, in which he was not offered a position to his satisfaction, and left for Canada, where he became involved in building up another AI research group. J.M. Foster, the other senior figure in the computer unit, went to Essex University, where R.A. Brooker, at that time chairman of the Computer Science Department, was in the process of building up AI research interests in an attempt to dispel the non-publishing lethargy prevalent there.[91] However, the situation at Aberdeen never became quite so fraught as at Edinburgh.

In Edinburgh, Michie attempted to fight back against the reorganisation. He marshalled support from the Dalle Molle Foundation to keep research going under his direction,[92] and he had previously published an article in the *New Scientist*, in response to Lighthill's report, asking why it was that AI, which required peanuts in financial terms, should suffer cuts, while nuclear physics, absorbing huge amounts of money and providing little proportionate return, should not be cut:[93] but to no avail. Essentially, by the time of the reorganisation he had lost the support of his colleagues, and Swann, who as Principal had supported and protected Michie, had left in 1973 to take up the chairmanship of the BBC.[94]

In the Machine Intelligence Research Unit after 1974, Michie's energies were channelled into promoting, directing and carrying out research on chess playing programs and organising further machine intelligence conferences, one in the United States with Elcock,[95] one in Russia,[96] and a third, again in the United States. He spent a considerable amount of time on visiting professorships abroad and took up scientific journalism, where in his regular column, 'Michie's Privateview' in *Computer Weekly*, he often commented on the importance of AI research. Latterly, in the late 1970s and early 1980s, he started vigorously promoting 'expert systems', (AI frameworks for representing and mobilising highly detailed specialist knowledge, which, given their essentially simple structure, have achieved remarkable levels of competence comparable to those of human experts), by organising conferences and schools to disseminate the approach to industry,[97] as well as directing and sponsoring relevant research in his unit.

Once the new regime of the Department of AI had settled down, activity did become normalised, with the conventional university emphasis on teaching coming to the fore. This owed much to the mid- to late-1970s financial stringencies which ensured that the universities looked to efficiency in their sphere of production — namely, the training of students. Undergraduate courses were experimented with at Edinburgh, and an AI textbook produced.[98] The application of AI in education itself — intelligent computer-aided instruction, clearly an area of direct social relevance — became a major concern at Edinburgh. This was an area originated by, among others, Seymour Papert, a colleague of Minsky at MIT, and was taken up by Michie and refined and developed at Edinburgh in a variety of approaches by J.A.M. Howe and his colleagues. In general, research activity remained at a high level, producing some 250 publications during the period 1975–80, and with a strong postgraduate school of about 30 being built up after an initial weakening due to the Lighthill report.

In some quarters, however, there was the impression that Edinburgh had declined in importance as an AI centre. This impression was partly due to the departure of highly respected researchers from Edinburgh, as already mentioned, but there were also other contributing factors. Firstly, Edinburgh no longer attracted the publicity over the robotics work as it had formerly done, though research with the robot equipment continued at a modest level; indeed, publicity was shunned as a matter of departmental policy, because of what was felt to have been over exposure by Lighthill. Secondly, an influential PhD thesis by T. Winograd at MIT in 1972[99] had brought the natural language research area to the centre of the AI stage, displacing the theorem proving approach, a shift which owed something to a 'witchhunt' against theorem provers led by Minsky.[100] Edinburgh, however, with a strong tradition in this latter area associated with the research group of Bernard Meltzer, at that time professor of computational logic and head of the department, had no strong competence in natural language and did not appoint a specialist on a permanent basis. Theorem proving had developed a very strong internal theoretical dynamic, deriving from mathematical logic and based on refining the 'resolution' method of machine-oriented inference devised by J.A. Robinson in 1965, which had brought it into the centre of attention in AI the first place.[101] This internal dynamic, coupled with the theorem provers' assured confidence in their formal mathematics-based status, had led to their being always rather autonomous; and under easier funding conditions, they would probably have become a completely independent specialty of computer science, much as pattern recognition, itself based on a strong internal dynamic, had done in the early 1960s.[102] However, the theorem provers made somewhat of a comeback in the late 1970s with the logic programming approach embodied in the language PROLOG, thus re-establishing themselves to some extent as a source of techniques of utility to AI in general.[103]

A third contributing factor to the perceived decline of Edinburgh as an AI centre was that computation theory, another of the major research themes in the department there, under the leadership of R.M. Burstall, professor of AI, had become more central to computer science in general during the 1970s and less of a specifically AI approach.[104] This situation was rationalised in the late 1970s with the transfer of the computation theory group from the Department of AI into the Department of Computing Science. A fourth and final contributory factor to the perceived decline of Edinburgh was that other major concentrations of AI interest had developed in Britain: at Sussex, Essex, and later in the Open University. At the same time, numerous one-person AI projects were pursued elsewhere, often based on 'colonisation' or 'infection' from the established centres.

Professor R.A. Brooker, always an enthusiast for AI (he was one of the panel members in favour of AI in the 1972 SRC *Computing Review*), had wanted to invigorate the research atmosphere in the computer science department at Essex, and to this end had recruited J.M. Foster from the Aberdeen AI research group to a chair in the department.[105] However, this did not work out and Foster left after about a year, having done little on the AI research side, but having developed the elements of a course in the area. There was considerable interest in the AI approach on the part of young researchers in the department such as J.M. Brady and R. Bornat, and when P.J. Hayes arrived in late 1972 from Edinburgh, bringing with him his extensive familiarity with the AI literature and research front, he catalysed the development of research projects in the field at Essex. One such project, supported by the SRC, was the development of a system to read handwritten FORTRAN coding sheets using high-level knowledge to guide the interpretation — a project of clear practical utility (although it never paid off) that nevertheless incorporated the AI approach.[106]

Once established as a centre for AI, other people were attracted there. J.E. Doran, who had been an early member of the Experimental Programming Unit and the Department of Machine Intelligence and Perception in Edinburgh, joined in 1973, bringing with him his research interests in using AI techniques for the reconstruction of cultural evolution in prehistoric settlements from data arising out of the archaeological excavation of graves — yet another illustration of the divergent nature of the AI activity.[107] Bruce Anderson, also from the Edinburgh AI centre, joined for a while before going to the Department of Electrical Engineering, while Yorick Wilks, who made a name for himself with his AI work on natural language and his numerous publications, was appointed to a chair in the Department of Linguistics — marking a new stage in the penetration of other subject areas by AI.[108]

This blossoming of AI interests at Essex was aided by the energetic and aggressive activity of Brady. He rapidly established himself as a competent researcher in the field and became involved at the organisational level, being elected to the chairmanship of the AISB society in the late 1970s, and organising a summer school which attracted leading researchers from the United States. He soon became known in the international AI community and made frequent visits to America, eventually leaving Essex to go to MIT in 1979. The MSc course and research environment at Essex proved to be an effective medium for the training of AI practitioners, some of whom subsequently went on to start AI groups elsewhere — notably H.J. Siekmann, who built up a group in Germany; B. Wielinga in the Netherlands; and C. Bearden, who founded an organisation similar to AISB in New South Wales in Australia. The level of interaction with other centres was high, particularly

those in Britain. There were links with Edinburgh through Hayes's connections there, and increasingly with Sussex through exchanges of students. In particular S. Hardy, after doing the MSc course at Essex, went on to do doctoral work in AI and then moved to Sussex, where he was largely responsible for the design of an AI computing environment for the cognitive studies programme.[109]

The emergence of Sussex as a major AI centre in the mid-1970s had a basis that went back to the mid-1960s. N.S. Sutherland, an experimental psychologist who had been interested in the cybernetic and information theory developments of the 1950s, was appointed to a chair at the newly-established University of Sussex. He was favourably disposed towards the AI approach, about which he had heard while he was at MIT in the early 1960s, and was much impressed by the ideas of M.B. Clowes, whom he had met at Oxford and who, as already mentioned, had founded the AISB society. Moreover, AI had been taught by Sutherland as an ingredient of the experimental psychology course from the start at Sussex, and he had arranged for two people from Edinburgh, Burstall and Doran, to visit Essex on a regular basis to give lectures on technical aspects in the area. As part of the conditions for Sutherland coming to Sussex, he had insisted on the founding of a brain research institute, and had negotiated with A.M. Uttley, Gregory, and Longuet-Higgins, for them to join him.[110] The latter two decided to go to Edinburgh instead, but Uttley came to Sussex and started work on the simulation of networks of an artificial neuron (the informon),[111] and on the application of these to perception. This was clearly a cybernetic rather than an AI project, but the computer provided for the project by the SRC made possible a subsequent characteristically AI attack on machine vision, which was carried out by Clowes when he returned from Australia in 1969 — work which rapidly established him as a leading figure in machine vision research.[112]

At that time in Sussex there was a broad base of interest in AI, favoured by the explicit focus on interdisciplinarity of the distinctive Sussex 'school' organisation which contrasted with conventional departmental divisions and their associated impervious boundaries,[113] while Asa Briggs, then vice chancellor of the university, was supportive of new ventures.[114]

In 1970 Sutherland put forward a radical proposal for a new School of Cognitive Studies with an intellectual focus on knowledge and understanding to include teaching and research in a range of subjects: computing science and AI; experimental psychology; linguistics; logic and philosophy; and mathematics.[115] However, in the restrictions on growth of the universities in the early 1970s, this proposal was turned down. However, a more modest development, which came to be called the 'Cognitive Studies Programme' was eventually launched within the existing School of Social Sciences, and in 1973 Clowes moved from experimental psychology to take up a chair in AI instituted for the programme. Several other members of the university were associated with the programme, among them M.A. Boden who had become familiar with the computational approach during her doctoral research in the United States on purposive behaviour in psychology,[116] and whose book *Artificial intelligence and natural man*[117] is one of the most accessible introductions to work in the area; and Aaron Sloman, a philosopher with a mathematics and physics background who was influenced by Clowes and Boden to consider work in AI and who spent a year in the Edinburgh centre before returning to start research in vision at Sussex.[118] His book, *The computer revolution in philosophy*[119] argues enthusiastically for the great potential of the AI approach in matters philosophical. Meanwhile, in 1974 Longuet-Higgins and some members of his research group from Edinburgh had joined the Centre for Research in Perception and Cognition, a research unit associated with the Laboratory of Experimental Psychology. Following Longuet-Higgins's move, other members of experimental psychology, Professor P.N. Johnson-Laird and C. Darwin, also developed research interests in AI,[120] which was taught as a compulsory part of the course presented there.

These two AI-oriented groupings ensured that Sussex emerged as a major centre for AI, especially in the areas of language studies and vision, in the 1970s. Moreover, the emphasis on the 'soft' applications of AI — the study of cognition, rather than on the 'hard' areas like the computer science and engineering applications — marked the rise to importance of cognitive science — the computational approach to linguistics, psychology, and philosophy, based primarily on the methods of AI. This rise was encouraged by the formation of a Cognitive Science panel in the SRC. Similar cognitive science concentrations emerged elsewhere, usually as a supradepartmental federation for research, as in the Edinburgh

School of Epistemics (founded by Longuet-Higgins and others in 1969, but only having really taken off in the late 1970s) and the Cognitive Science Institute at Essex, also implemented at the end of the 1970s. At the Open University similar developments had occurred. The Cognitive Psychology course included a substantial AI component, while the presence of other researchers, interested in the application of AI to education, meant that the Open University itself comprised an emerging centre for AI in the late 1970s.

As well as being an essential element in the emergence of cognitive science, AI has become a recognised specialty within computer science.[121] While these developments in the 1970s bear some resemblance to Lighthill's predicted fission of AI research between the categories of computer-based central nervous system research on the one hand, and advanced automation on the other, it is difficult to align his prediction with the continued coherence of AI — that is, the continued identifiable existence of the AI paradigmatic structure. Rather than the established areas of linguistics, psychology, and philosophy absorbing AI, it would seem that cognitive science is an emerging synthesis based on the unifying computational modelling approach of AI. Indeed, one could equally well argue, against Lighthill, the AI practitioner's extreme view that what is happening is merely the process of colonisation of other areas by the AI approach: 'I see the future of AI as a very long haul towards computational theories of physics, chemistry, linguistics, sociology, visual perception, locomotion and every other aspect of what it means to be human'.[122] Clearly, therefore, views still differ over the assessment of the place and future of AI.

3.7 THE ESTABLISHMENT IN THE UNITED KINGDOM

One of the main features of the development of AI in Britain was the initial and continuing strong American influence. Nearly every one of the leaders of AI research in Britain had visited the United States and been impressed by developments there, before moving into the area, or promoting it themselves in Britain: Michie, Meltzer, Sutherland, and Clowes had all visited AI projects in the United States in the early 1960s. These links with the American AI community were maintained and strengthened in the ensuing years, and continue today with a high international exchange of personnel between the various centres.

These links also underlie the substantive similarity of research pursued in Britain and America, at the general level of the paradigmatic structure, as reflected for instance in the research profile at Edinburgh which matched the patterns evident in the United States. This has remained true since the mid-1960s when AI took off in Britain, and it has largely been the case that the initiative in the development of AI in terms of the broad content of research has remained in the hands of the establishment in the United States. Consequently, the emergence of the establishment in Britain has been bound up with organisational aspects of development to a greater extent than in the United States, where organisational and general substantive innovation were both important. Nevertheless, just as in America, the emergence of the establishment in Britain has been inextricably bound up with the development of the field. The organisational aspects of development have not only predominated in Britain but have also had an appreciable impact on the shaping of the wider international AI community: for instance, the Machine Intelligence Workshops played an important role leading to the establishment of a journal and an international conference structure for the area; the first dedicated AI department was established at Edinburgh; and the cognitive science concentration has received its firmest institutional expression in Britain, with the Cognitive Studies Programme at Sussex, and the School of Epistemics at Edinburgh.

Moreover, in the development in Britain, a clear division between what have been called organisational and intellectual leadership roles[123] is identifiable, especially with respect to the first generation establishment. At Sussex, for instance, N.S. Sutherland was energetic in supporting the AI approach there, getting grants for researchers to work in the area without himself being actively involved. Similarly, at Essex, R.A. Brooker deliberately encouraged the development of an AI research group within the Department of Computer Science, but did not himself actively contribute to research to any great extent. It was also the case at Edinburgh with Michie, whose contributions in substantive terms were overshadowed by his role as an organisational leader: indeed, he ranks as a scientific entrepreneur of the first order. He alone, almost single-handedly, was responsible for getting AI launched in the United Kingdom, and his influence appears to have lurked behind nearly every event of major importance concerned with AI in Britain in the 1960s.[124] While his enthusiastic promotion of robotics was eventually to backfire with Lighthill's condemnation, his activity was nevertheless instrumental in putting Edinburgh on the map with respect to AI, and for providing an environment in which young researchers were able to establish international reputations. Despite his fall from favour after the Lighthill report, Michie remained active — organising a further three Machine Intelligence Workshops; promoting and directing research on chess-playing programs; and, latterly, working to bring the 'expert systems' applications area of AI to the notice of a wider audience, including industry. Also at Edinburgh, Meltzer established himself as an 'elder statesman' of AI, channelling his energies to maintaining high critical standards in the specialty Journal, and to providing an environment in which major developments in theorem proving were made; while his own substantive contributions have been in the form of synthesising reviews, directing attention to certain problems and suggesting fruitful possibilities. Even Longuet-Higgins, with his dislike for organisational involvement, preferring purely scientific activity, has played a role of organisational importance, in arguing for the cognitive science applications of AI, and in helping to found the School of Epistemics at Edinburgh.

An important part played by the organisational role has been the provision of facilities and opportunities for young researchers and students to rise into the establishment, in many cases developing an international reputation for themselves with their PhD work. The high number of important contributions to the field at PhD level seems to have been characteristic of AI, and has combined with the very fluid and informal nature of the organisation of work in the area, to maintain a shallow internal hierarchy with no elaborate vertical division of labour. Terry Shinn's description of the organisational structure in laboratories concerned with computer research in vector analysis could apply equally well to the case of AI.[125] While this organisational fluidity could stem to some extent from the holistic and diffuse nature of the research goal of modelling human intelligence, it seems likely that the situation also arises out of the youth of the field. The preponderance of young researchers at about the same stage in their professional careers, and hence of more or less equal status, together with the general shortage of personnel with the appropriate computer background, has rendered it difficult to support a steeply hierarchical structure — even where such a structure *would* be possible, as, for example, in some well routinised commercial computing where there has already been a drive towards extreme stratification and division of labour along scientific management lines.[126] However, there has been a horizontal division of labour, with specialisation in different research areas and, of course, organisational roles have tended to be filled by professionally senior practitioners. It is likely that the taking up of such roles by the first-generation establishment was the outcome of two factors: on the one hand, the clear opportunities offered by these roles and their availability in the early days of the emergence of AI in Britain; and on the other hand, the difficulty of getting to the research front without extensive specialist training, particularly in the activity of programming, at that time very much an esoteric art. Moreover such training obviously required some sort of organisational structure for efficient transmission. Consequently, the organisational roles were more important and probably demanded more effort during the initial emergence of AI in Britain, before they became institutionalised and suitable recruits were readily available from among the ranks of the trained AI specialists. With the development of the institutional structure during the 1960s and 1970s, there has emerged a similar tight pattern of intergenerational and intercentre linkages in Britain as that evident in the United States: indeed, the American and British structures were connected, as is illustrated by Figure 3.2.

This tight intergenerational pattern has been reinforced by two related characteristics of the AI paradigmatic structure: the *constructive* and the *craft* nature of work in the area. The writing of a computer program to carry out some task clearly involves making or constructing something, rather than investigating something that is naturally given (although such investigations can be and are involved). Not only are there many ways of constructing a program on AI principles, but there are also many other non-AI ways — for instance, the construction of a stochastic model. Furthermore, it is possible to construct models not based on the use of the digital computer — for instance, in the building of electronic analogs of neuron

Figure 3.2: Movement of research personnel between AI centres in the United Kingdom, 1964–80. Each line shows the trajectory of an individual; ●—— represents the start of a career in AI. The density of horizontal lines gives some indication of the frequency of movement between centres. The figure represents trajectories for 50 out of a total of about 150 personnel involved in AI research in Britain during the period 1964–80; a further 43 did not move, and 15 left the field altogether. Information was not available in the remaining 40–50 cases, and is less complete after 1978, while short-term visits of six months or less (which were quite frequent) have not been included. Thus, the above is probably an understatement of movement between centres. The following features are apparent: strong links with the United States; movement away from Edinburgh in the mid-1970s; increased movement in the early to mid-1970s involving the newly emerged centres at Sussex and Essex; increasing links with centres elsewhere in the 1970s; and finally, the importance of Edinburgh, and latterly Essex, as sources of AI personnel.

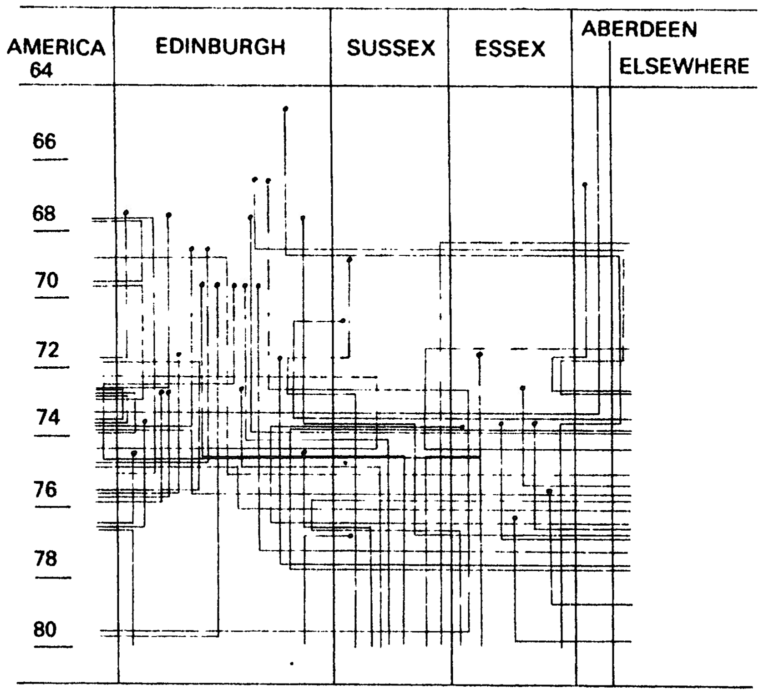

networks. Hence, it is clear that the AI paradigmatic structure delineates only one broad way of making models among many possible alternatives; this arises from the constructive nature of the activity in the area. Central to the paradigmatic structure is the activity of programming, based on the use of list processing languages, associated with which are many characteristic techniques and 'tricks of the trade'. Effective programming also involves a high degree of skill, and all of these features together lend the AI approach a distinctly 'craft' nature, which requires for its transmission a lengthy apprenticeship and some degree of contact and interaction with experienced practitioners. This craft and constructive nature of work in the area ensures that it is extremely unlikely that the specifically AI approach, even in broad terms, would be developed spontaneously and independently outside the community of people already using it. Consequently these aspects of the paradigmatic structure place constraints on the development and transmission of the AI approach, and thus reinforce the tight intergenerational pattern, which is also encouraged by the purely social aspects of communication and the favouring of those already known in the network.

A major feature of development in Britain was that the initial AI establishment did not emerge from the strong pre-existing cybernetic or computer science network, despite the clear prefigurement of AI research in the work of Turing and Craik. Rather, it came from people external to such work: Elcock at Aberdeen; Michie, Meltzer, and Longuet-Higgins at Edinburgh; and later, Boden and Sloman at Sussex. Even Sutherland and Clowes at Sussex, although they had links with the cybernetic tradition, were certainly not centrally involved with it. Such an entry into the AI establishment from outside the cybernetics tradition had also been evident in the United States, with Simon and Newell. However, just as was the case in the United States, the first generation members of the British AI establishment, although they were marginal to cybernetics and computer science, certainly did not arrive from nowhere. Michie, for instance, had some 60 publications to his credit in his previous specialist areas of genetics, immunology, and reproduction, and held the post of reader in the Department of Surgical Science at Edinburgh. He was recommended by W.H. Waddington and M. Swann on his appointment to Edinburgh in 1958, and Swann continued to support him during his AI activities. Meltzer, as a

reader in the Department of Electrical Engineering, also had a solid reputation for his work on electron beam dynamics (used by NASA for the design of ion propulsion for space vehicles) and solid-state electronics; and Longuet-Higgins, with his international standing in theoretical chemistry was clearly already a member of the wider British scientific establishment.

The migration of outsiders with some standing, and hence the freedom to move fields,[127] therefore seems to have had as its major consequence the construction of an organisational structure within which the pursuit of AI research subsequently developed. In some cases (especially that of Sutherland) the organisational leadership role merged with another, probably necessary, role in the emergence of a new interdisciplinary area — namely, a sponsorship role. The institutional developments at Edinburgh would not have been possible without some strong support from sponsors placed in fairly influential positions with the university government, especially in view of the tendency for the publicity attracting, research intensive, AI activities to arouse suspicion and resentment: both Sir Edward Appleton and Sir Michael Swann, who succeeded Appleton as Vice Chancellor of Edinburgh University, were active in encouraging and supporting those developments; and at Sussex, there was also fairly widespread support for AI-oriented developments among many of those in positions of influence.

As well as this positive sponsorship within the wider establishment, which was accompanied by a positive evaluation of the status of AI, there was a negative sponsorship as well, as demonstrated by Sir James Lighthill's report. Lighthill's opposition could clearly be seen to support the *status quo*: he affirmed the value of the currently existing areas included in advanced automation (e.g. control engineering) and in central nervous system research (e.g. neurophysiology), and moreover he attributed any success in AI to contributions arising from these areas. The validity of the emerging interdisciplinary area of AI was thus challenged and denied, explicitly in terms of the already-established disciplines surrounding AI. In particular, Lighthill's focus on the established central nervous system areas of research — neurophysiology and neurochemistry — and his use of the term 'central nervous system' happened to align with the dominance in Europe of the neurosciences which study the 'hardware' of the brain, over the cognitive sciences — linguistics and psychology — which might be said to study the 'software' of the brain, a dominance which is not as clear-cut in the United States.[128] In this dominance relation we see a prestige hierarchy, with those sciences closest to the physical sciences accorded most prestige. In this context, the study of the brain at the reductive level of biochemical or neurophysiological mechanisms is considered more prestigious than the AI approach at the information-processing levels.

Moreover, it is difficult to account for the impact of the Lighthill report (by his own admission a two-months layman's view of the area)[129] except in terms of the authority carried by Lighthill's eminence. It is interesting to note that reference is still offered to observations that Lighthill made on the difficulties for search arising out of the combinatorial explosion,[130] as if he were their originator, whereas, in fact, the combinatorial problems had been regarded as the *raison d'être* for AI — the huge size of the space of possible moves in chess, for example, estimated by Shannon in 1950 as in excess of 10^{120}, dictated the need for heuristic strategies to restrict the search space to manageable proportions.[131]

What is also of interest about the Lighthill report, however, apart from its importance as an authoritative pronouncement on the status of AI research, is that it is one among a multitude of attacks on the field.[132] Such attacks on AI have been commonplace, and while they purport to deal with the particularities of the subject matter of research in the area, it is quite clear on closer inspection, that they are more concerned with the general goal of constructing an intelligent machine. It would be too lengthy to argue this fully here, but it is perhaps pointed by the continued relevance of Turing's comments on the arguments for and against the possibilities of constructing intelligent machines,[133] despite the fact that the distinctive AI approach had not emerged when Turing was still alive. It is also pointed by the fact that these attacks on AI are not the prerogative of any particular group: criticism has come from all shades of political opinion, and from all areas of research, scholarly as well as technical.

Ironically, this variety in attacks can only be matched by the diversity in the sources of support for AI, or the range of (often conflicting) views within the field itself. Something of this has already become evident in the differences between Longuet-Higgins and Michie, one favouring the cognitive science definition of AI, and the other the machine intelligence

approach. There are other divisions: those supporting a theoretical formal approach, such as McCarthy, for instance, and those supporting the exploitation of practical applications, such as Feigenbaum; those who see no problems with accepting military funding (McCarthy and Feigenbaum) and those implacably opposed (Meltzer, while Michie was opposed to classified work).

These divisions, which are legion in the area, are coupled with a rather amazing state of substantive partisanship, or scientific ethnocentricity, in which proponents of the various different research areas each tend to see their own approach as the real AI approach — to the theorem provers, theorem proving is central; for the natural language proponents, language is the basis for reason, and so on.[134]

Many of these differences can be related to the background competences of the practitioners, and can be interpreted as competition between groups on the research area level for resources and authority within AI as a specialty, constituting perhaps the primary locus for competition over cognitive commitments, quite distinct from the individualistic level of competition for recognition, long identified in the sociology of science as a basic motor of scientific development. An important point to note is that such differences are not precluded by the paradigmatic structure of AI, outlined earlier as providing guidelines for the common computational approach and its programming basis, but which does not dictate a dogmatic monolithic attitude, nor inculcate a unifying solidarity. Moreover, this variety of views in and around AI can be related to its position as an interdisciplinary area, with particular research areas associated with particular neighbouring disciplines — for example, the natural language research area is associated with linguistics, while theorem proving has links with metamathematics; and as an interdisciplinary area, the status of AI research is still very much in process of negotiation. The cognitive science developments appear to have led to an acceptance, on the part of those involved of the validity of AI: indeed, the impression in that context is that AI is the 'hard' formal core, and therefore of high status. However, in a computer science or general scientific context, AI is still seen very much as a 'freaky', rather dubious fringe activity, and consequently of rather inferior status.[135] Moreover, there are two broad categories of attacks which can be related to these contexts. On the one hand there are attacks, often by philosophers and others, on AI for being reductionistic and impossible[136] — in a sense it is 'harder' than is appropriate for the study of intelligent activity. On the other hand there are criticisms, often by computer scientists, of AI for being morally wrong,[137] bad science,[138] or undisciplined and sloppy.[139]

Finally, perhaps the variety and depth of feeling of the many attacks on AI derive not so much from what is in fact done in AI research, but rather from the fact that the very broad aim of research in the area — namely, the construction of intelligent machines — bears uncomfortably on our conception of ourselves. In Elias's terms, AI research is seen to be involved with a very sensitive area of the means of orientation: the area which is concerned with the nature of mind. Furthermore, during several hundred years of development and struggle with other competing establishments, a scientific establishment has yet to succeed in gaining a monopoly over the means of orientation in this area. In making its challenge in this area, therefore, AI is inviting violent attacks, and its practitioners should hardly be surprised when they suffer them.

3.8 CONCLUSION: THE PROCESS OF DEVELOPMENT AND ESTABLISHMENT IN ARTIFICIAL INTELLIGENCE

The process of development and establishment in AI in broad terms therefore appears to have been as follows. Around the period of the Second World War, catalysed by the war-time weakening of traditional disciplinary boundaries, and brought into being by exigencies deriving from the unprecedented problems of organisation and communication posed by the increasingly complex social structures and conditions, there emerged the software sciences. These had their focus on pattern rather than substance, and included operations research, computer science, and cybernetics. In particular, the cybernetics area, with its focus on the processes common to animals and machines, promised a realisation of the age-old desire to make an artificial human, a machine that could think. Within the general area of cybernetics, various approaches were made to the construction of intelligent machines, some based on electronic analog of neuron networks, others on the simulation of processes by means of the newly developed digital computer.

In this context, the paradigmatic structure of AI was articulated in the United States during the late 1950s.

The American establishment in AI consisted essentially of those who had contributed to this articulation, and who had provided the institutional structure within which subsequent research, based on the distinctive AI paradigmatic structure, could be undertaken. The source of authority and reputation of the establishment was derived from the effective demand which developed for this distinctive approach on the part of those who wanted to follow the approach themselves and those who thought the approach was worthwhile and promising. In addition, the emerging establishment secured the backing of the funding agencies, aided by their good connections with the wider establishment, in competition with other approaches within computer science. The preference of the funding agencies for concentrating resources in a few centres, coupled with the expense of the instrumental base required (the digital computer), ensured that an effective monopoly over material resources as well as the cognitive ones was maintained.

As the paradigmatic structure was already elaborated before it was exported to the United Kingdom in the early to mid-1960s, it left less opportunity for substantive contributions on this general level while there was ample scope for extending the institutional facilities for carrying on AI research in Britain: consequently, the first-generation British establishment in AI consisted of those who were able to set up organisational forms to exploit the already articulated paradigmatic structure, thus giving them a monopoly over the cognitive resources in the area. Because of the constructive nature of AI, which was manifested in it being only one among several competing cybernetic approaches, and since people already committed to a particular approach tend to stay with that approach, members of the first-generation establishment included people from outside the cybernetic and even computer science traditions. This was the case in the United States, but was more marked in Britain. Moreover, the members of the first-generation establishment were, in fact, drawn from those who already had some standing or prestigious backing in another field; and being thus well connected, they were able to secure the backing of the Science Research Council in competition with other approaches within computer science, ensuring their effective monopoly over material as well as cognitive resources in the area.

This monopoly over resources was reinforced by the constructive and craft nature of the AI paradigmatic structure, and the patterns of development in both the United States and the United Kingdom followed the lines of personnel mobility and contact, thus leading to a tight intergenerational and intercentre structure of linkages. In particular, students and 'descendents' of the members of the first-generation American establishment have tended to dominate the field. The paradigmatic structure of AI, however, while providing general guidelines for the methodological approach employed in AI research, does not dictate the direction of research. Due to the very wide-ranging nature of the focus on intelligent activity, research in AI has become involved with the subject matter of many other disciplines, and has therefore developed as an interdisciplinary area. This interdisciplinary character of AI has induced many mutually competing divisions *within* the area, as well as leading to many external views on the status of the field, which has consequently been very much a matter for negotiation: from the point of view of the 'soft' sciences, such as linguistics and psychology, AI has appeared 'hard' and therefore of superior status; from the point of view of the 'hard' sciences, such as computer science and physics, AI has appeared somewhat 'freaky' and therefore of inferior status. Furthermore, as a newly emerging specialty, AI has been in competition with already established disciplines: the Lighthill report, critical of AI, can be interpreted in this context as an affirmation of the *status quo*. Finally, because of the general aim of constructing intelligent mechanisms, AI has been seen as challenging the monopoly on the means of orientation with respect to the nature of mind. As this bears directly on peoples' conceptions of themselves, deep feelings have been aroused, as is evident in the many and varied attacks on AI.

It is therefore evident that there has been competition on a variety of levels over AI, with differing consequences for the development of the field, and, moreover, engaging distinct groups or establishments. At the most circumscribed level, *within* the field, there has been competition over the choice of techniques to be used in a particular research area, and competition between research areas themselves. This is of immediate consequence for research at the practical level, and involves groups negotiating for ascendency within the AI establishment. At a less circumscribed level, the establishment

of AI as a whole has been involved in competition with neighbouring scientific and scholarly establishments. This concerns the general scientific validity of the field rather than being of immediate practical import, and has consequences in broad terms for the status of the field, the availability of funding, and the continued demand for the AI approach on the part of potential entrants. Finally, at the most general level, and most clearly underlining Elias's remarks about the means of orientation, the validity of the AI approach is discussed in terms involving political, religious, moral, philosophical, and cultural issues. The question being asked here is not whether AI is valid as a scientific approach, but rather whether any such approach to the mind is viable. This engages a far wider group, and the prevalence of the debate at this level is perhaps pointed by the fact that, at most, a couple of dozen full-time occupational opportunities are available in AI — very few compared with many other areas of endeavour — while nearly everyone, it seems, has something to say on the issue of whether machines can think.

3.9 POSTSCRIPT: THE COMMERCIALISATION OF ARTIFICIAL INTELLIGENCE

Since the foregoing was written in 1980, there have been dramatic developments in the general reception afforded to AI, with large increases in public funding for research in the area, and even larger increases in commercial interest and financing. Indeed, the early 1980s marked the transition from a phase of 'establishment' to one of 'commercialisation', in which attempts were made to convert or 'transfer' (the fashionable term) the results of AI research from the realm of the merely fascinating into commercially useful products which could be sold, it was hoped, for a profit.

This recent phase of development was triggered by the news that the Japanese Government planned a large-scale programme to put Japanese industry into the forefront of world computer developments, an area where, unlike many others such as manufacturing, the Western nations still had a significant lead.[140] The *Summer 1980 JIPDEC Report*, (JIPDEC: Japan Information Processing Development Centre) put it bluntly:

> . . . the Ministry of International Trade and Industry initiated its Fifth-Generation Computer Development Project, a project which, with a target year of 1990, has the aim of developing truly world leading computer system technology and promoting the development of Japan's computer industry through research and development into fifth-generation machines.[141]

This report was published (in English) in September 1980, and outlined the Japanese plans, emphasising the need for supporting and exploiting the results of AI research, especially natural language and knowledge engineering work. These plans were heralded with great excitement, even trepidation: for instance, in the United Kingdom, Alex d'Agapeyeff, an eminent computer software entrepreneur,[142] and at one time a president of the British Computer Society, claimed at a meeting in London in December 1981 that the Japanese were declaring 'economic war' on the West; and that the language they used was that of *Mein Kampf*.[143] D'Agapeyeff went on to play a further part in mobilising concern, by conducting surveys and organising other initiatives.[144]

There were similar reactions elsewhere. In the United States, Edward Feigenbaum, one of the leading 2nd generation American AI practitioners,[145] and Pamela McCorduck[146] collaborated to produce a book about these developments, which called for a concerted American response:

> America needs a national plan of action, a kind of space shuttle program for the knowledge systems of the future. In this book we have tried to explain this new knowledge technology, its roots in American and British research, and the Japanese Fifth Generation plan for extending and commercialising it. We have also outlined America's weak, almost nonexistent response to this remarkable Japanese challenge. The stakes are high. In the trade wars, this may be the crucial challenge. Will we rise to it? If not, we may consign our nation to the role of the first great postindustrial agrarian society.[147]

This book became a best seller, selling 10,000 copies within a matter of months in the Japanese translation alone.

The Japanese Ministry of International Trade and Industry

(MITI) announced the commencement of their ten-year, US$850 million effort in October 1981, and other national governments were not long in responding with comparable initiatives. In the United Kingdom, the Alvey programme was started in 1983, following the recommendations put forward in a 1982 report, *A programme for advanced information technology*,[148] by a study committee set up as a more or less direct response to the Japanese plans. Funding of some £352 million was proposed, of which £26 million was to be spent directly on Intelligent Knowledge Based Systems (essentially relevant AI research), with more for AI through the £78 million funds for 'demonstrators' and education. As Patrick Jenkin, the Secretary of State for Industry, said, in a statement to the British House of Commons:

> This is the first time in our history that we shall be embarking on a collaborative research project on anything like this scale. Industry, academic researchers and Government will be coming together to achieve major advances in technology which none could achieve on their own. The involvement of industry will ensure that the results as they emerge are fully exploited here in Britain to the advantage of our economy. Information technology is one of the most important industries of the future and therefore one upon which hundreds of thousands of jobs in the future will depend.[149]

In the United States there were several comparable programmes, including the Microelectronics and Computer Technology Corporation (MCC), and the Defense Advanced Research Projects Agency $1 billion Strategic Computing programme. The former was officially launched in January 1983, after nearly a year of meetings, while the decision to go ahead seriously with the latter was made in spring of 1982,[150] following the news of the Japanese initiative. Many other relevant programmes have also now started around the world, notably the multi-nation European Strategic Programme for Research and Development in Information Technologies (ESPRIT), and the controversial American Strategic Defense Initiative (SDI), popularly known as 'Star Wars'.

Clearly, therefore, the climate of reception for AI in the 1980s changed considerably from the sceptical years of the Lighthill report. Catalysed by the innovative example of the

Japanese intiative, information technology was placed firmly on the agenda of every major government, with AI at the centre of attention and effort.

In the UK, this resulted in a large growth in the numbers of people interested in AI, and in employment opportunities in the area. At the end of the 1970s, there were perhaps a dozen or so full-time practitioners, with approximately another hundred people with significant interest; by 1986 there were probably over one hundred full-time AI posts (though not necessarily occupied by fully trained or experienced AI practitioners!), over one thousand people with significant interests, and many more sufficiently interested to attend the increasingly expensive conferences on the subject.[151] From the handful of UK academic institutions involved in the 1970s, by the mid-1980s nearly every institution of higher education in the UK showed signs of serious interest in AI research and training.[152] By the mid-1980s also, substantial industrial and commercial involvement was developing, with many companies seriously looking at the potential for making use of AI systems in their own operations.[153] An AI supply and service infrastructure was emerging from what had been a very small basis indeed in the 1970s, with perhaps three or four companies supplying AI products of some sort: by 1986, there were some 24 new start-ups (and at least two close-downs), and in addition, many major established electronics and computer firms were developing in-house AI divisions.[154]

The basic patterns of research in AI outlined earlier (in section 2 of this chapter) and in existence by the early 1960s, as discussed in section 4, were still essentially recognisable in the late 1980s. Indeed, it became routine to talk of AI 'tools'[155] while the research area structure continued to develop and differentiate. There were, however, perhaps three major changes affecting the development of the substantive structure of the area, which came to the fore during the 1980s.

The most obvious was in the rise of the so-called 'expert systems' area, with its emphasis on 'knowledge engineering' or knowledge-based information processing. This was seen to be the leading edge as far as practical commercial exploitation was concerned, and attempts were made at building expert systems to deal with a wide range of practical domains.[156] It was also the area identified by the Japanese as being at the core of the 5th generation computer systems.[157] Essentially what happens in

the development of an expert system is that AI knowledge-representation methods and techniques are used to capture the knowledge of a human expert in the target domain, to produce a knowledge base. This base is in turn processed by other AI techniques for making inferences, to draw appropriate conclusions as required, and thus emulate the intelligent reasoning of the original human expert. The development process for expert systems requires close interaction between the AI practitioner or 'knowledge engineer', and the domain expert.

Another major change, although not one with immediate or direct implications for the commercialisation of AI, was the rise of Cognitive Science. This referred to the use of the computational model in the human sciences, often in a purely metaphorical sense rather than necessarily implying the implementation of a computer program, and it has been argued that this represents a paradigmatic shift for the human sciences.[158] Given the massive attention afforded AI through the developments outlined above, the exploitation of computational ideas in this way was hardly surprising, and compares with the similar exploitation of conceptual resources arising from previous 'new' technologies such as hydraulics or clockwork mechanisms.[159]

The third major change affecting the structure of AI, and one which in fact underpinned and supported the other two changes, as well as underlying the broader field of information technology in general, was the dramatic decrease in computing costs. This lowered the hardware resource barrier to carrying out AI research, a barrier which had been very much in evidence during the early years of development, as previously discussed. As a result, the scope for a strong AI establishment to dictate and control development in the area was considerably weakened, allowing outsiders to move in and pursue their own variants of AI research.[160] Low computing costs should also encourage the rapid and wide diffusion and use of expert systems and other AI products, once their viability and efficacy has been demonstrated. Indeed, it became possible in the early 1980s to buy an expert system 'shell' (that is, the basic structure of an expert system, without the domain specific elements) for a personal computer, although the general efficacy of such systems appears rather constrained.

However, despite the highly favourable climate and the dramatic growth of effort in AI, debates over the validity of the approach and associated claims remained as lively as ever.[161]

Indeed, it could be argued that the far bigger market of interest in AI provided an even better basis for making a reputation for oneself by producing a critique of the area.[162] AI was still as popular as ever a focus for attack: clearly the struggle for monopolisation of the means of orientation continued unabated.

It was also not clear, at least by the mid-1980s, that practical commercial success had been consolidated or even proven, whatever the promise of research or prototypes. Some systems, in actuality constructed along more or less conventional lines, have been somewhat misleadingly described as expert systems, while in any case opinions differ widely over what really constitutes an expert system, and over just how many are in existence.[163] Even firms in the emerging supply and service infrastructure appear to make most of their sales to other suppliers and research and development teams, rather than to a wide base of satisfied users.[164] Furthermore, despite the increases in funding and personnel, it was also not clear that research results had benefited, at least not by the mid-1980s. A survey for one of the Alvey monitoring efforts (set up to monitor progress on the Alvey programme), which looked at publications in AI, found that while the total number of items published increased dramatically, this increase was almost totally accounted for by review articles rather than by more substantial pieces of work.[165]

However, this situation should not be surprising, in the light of the identification of the tight intergenerational craft-constrained nature of development in the area, as discussed earlier in the chapter. At the core of this process of development is a high degree of what might be called 'tacicity', by which I mean that the craft knowledge and practices which constitute the core burden of AI expertise have not yet been made sufficiently explicit, nor sufficiently generalised from the contingencies of their development, to be readily transferable without extensive hands-on experience or appreciable apprenticeship periods: in short, they are still largely tacit. If this tacicity hypothesis is valid, then, given the very small numbers of fully experienced or trained personnel in existence at the end of the 1970s, and given the demands on those people to carry out all sorts of tasks such as research and industrial consultancy as well as training, it is clear that significant effort will be slow to build. Meanwhile, most of the effort by the community as a whole needs to be taken up with learning;[166] and what better way of learning about a new area than writing reviews?

The question nevertheless remains whether AI is in fact yet sufficiently mature for such whole-scale exploitation, as it is in scientific terms a very young field, and, moreover, one which is attempting to tackle extremely difficult and profound problems. The doubt will surely linger that, while the hot-house climate of commercial exploitation and abundant fertiliser of industrial funding will bring on certain exotic fruits, others, perhaps more subtle or sensitive, and possibly more rewarding in the long term, will suffer.

ACKNOWLEDGEMENTS

I should like to acknowledge the help given by the interviewees to whom a draft of this paper was circulated. The comments of J.A.M. Howe, H.C. Longuet-Higgins, D. Michie and N.S. Sutherland were particularly valuable. I should also like to thank R.D. Whitley for his suggestions and the SRC for its support. Finally, thanks are due to the ESRC for support relevant to the writing of Section 3.9, the postscript.

NOTES AND REFERENCES

1. The term specialty is used here in the sense of a community of practitioners differentially located with respect to a common paradigmatic structure which defines a general focus of attention. Compare R. Whitley, 'Components of scientific activities, their characteristics and institutionalisation in specialities and research areas' in K. Knorr, H. Strasser, and H.G. Zilian (eds), *Determinants and controls of scientific development* (Reidel, Dordrecht, 1975), pp. 37–73.
2. See, for example, I. Aleksander, 'Artificial Intelligence', *Electronics and Power*, vol. 22 (1976), pp. 242–4.
3. This paper is based on a detailed MSc/PhD study of the area: J. Fleck, 'The structure and development of artificial intelligence: a case study in the sociology of science', unpub. MSc diss., University of Manchester, 1978. Also PhD diss. in preparation.
4. Compare R. Whitley, 'Components of scientific activities' (fn. 1).
5. Sandewall has discussed these aspects, for example, E. Sandewall, 'Programming in an interactive environment: the "LISP" experience', *ACM Computing Surveys*, vol. 10 (1978), pp. 35–71.
6. Compare J. Law, 'The development of specialties in science: the case of X-ray protein crystallography', *Science Studies*, vol. 3 (1973), pp. 275–303.

7. This is clear from a consideration of the contents of the book, a classic in the area, E.A. Feigenbaum and J. Feldman (eds), *Computers and thought*, (McGraw Hill, New York, 1963).

8. Compare D.E. Chubin and T. Connolly, 'Research trails and science policies: the shaping of scientific work by hierarchies and elites', in N. Elias, H. Martins and R. Whitley (eds), *Scientific establishments and hierarchies*, Sociology of the sciences, vol. VI (Reidel, Dordrecht, 1982).

9. D.O. Edge and M.J. Mulkay, *Astronomy transformed: the emergence of radio astronomy in Britain*. (Wiley-Interscience, New York, 1976).

10. Though for a dissenting view, see B.R. Martin, 'Radio astronomy revisited: a reassessment of the role of competition and conflict in the development of radio astronomy', *Sociological Review*, vol. 26 (1978), pp. 27–55.

11. See E. Yoxen, 'Giving life a new meaning: the rise of the molecular biology establishment', in N. Elias, H. Martins and R. Whitley (eds), *Scientific establishments and hierarchies*, Sociology of the sciences, vol. VI (Reidel, Dordrecht, 1982).

12. N. Elias, 'Scientific establishments', theme paper for Conference on Scientific Establishments and Hierarchies, Oxford, July 1980.

13. N. Wiener, *Cybernetics — control and communication in the animal and machine*, (Wiley, New York, 1948).

14. For instance, W. Pitts and W.S. McCulloch, 'How we know universals', *Bull. Maths. Biophysics*, vol. 9 (1947), pp. 127–47; or O.G. Selfridge, Pandemonium: a paradigm for learning' in *Mechanisation of thought processes* (HMSO, London, 1960), pp. 513–26.

15. Machine Translation was an early example of this approach. Based on syntactical analysis and dictionary look-up, it failed at that time in its aim of providing high quality translation — see Y. Bar-Hillel, 'The present status of automatic translation of languages', in F.L. Alt (ed.), *Advances in computers*, vol. 1 (Academic Press, New York, 1969), pp. 92–163.

16. C. Shannon and J. McCarthy (eds), 'Automata studies', *Annals of Mathematics Studies, No. 34* (Princeton University Press, Princeton, NJ, 1956).

17. Comments by J. McCarthy, at an 'AISB Summer School on Expert Systems', held at Edinburgh University, July 1979.

18. The Dartmouth Conference is mentioned in several places, for example, M.L. Minsky, 'Artificial Intelligence' in *Information* (a Scientific American book) (Freeman, San Francisco, 1966), p. 194; and in P. McCorduck, *Machines who think* (Freeman, San Francisco, 1979), pp. 93–114.

19. A. Newell and H.A. Simon, 'The logic theory machine', *IRE Trans. on Info. Theory*, no. IT-2 (1956), pp. 61–79.

20. G.W. Ernst and A. Newell, GPS: *A case study in generality and problem solving* (Academic Press, New York, 1969).

21. J. McCarthy, 'Recursive functions of symbolic expressions and their computation by machine', part 1, *Comm. ACM*, vol. 3 (1960), pp. 184–95.

22. M.L. Minsky, 'Some methods of Artificial Intelligence and heuristic programming' in *Mechanisation of thought processes* (HMSO, London, 1969), pp. 3–28.

23. In addition, the 'Matthew Effect' would be operating: R.K. Merton, 'The Matthew Effect in science', *Science*, vol. 159 (3810) (5 Jan. 1968), pp. 56–63.

24. J. McCarthy, 'Toward a mathematical theory of computation' in *Proc. IFIP Congress 1962* (North Holland, Amsterdam, 1963).

25. M.L. Minsky (ed.), *Semantic information processing*, (MIT Press, Cambridge Mass., 1968).

26. M.L. Minsky, 'A framework for representing knowledge' in P. Winston (ed.), *The psychology of computer vision* (McGraw Hill, New York, 1975), pp. 211–77.

27. See, for example, M.L. Minsky and S. Papert (and staff), 'Proposal to ARPA for research on Artificial Intelligence at MIT 1971–1972', *MIT AI Lab. Memo.*, No. 245 (Oct. 1971).

28. A. Newell and H.A. Simon, *Human problem solving* (Prentice-Hall, Englewood Cliffs, N.J., 1972).

29. W.S. McCulloch and W. Pitts, 'A logical calculus of the ideas immanent in nervous activity', *Bull. Maths. Biophysics*, vol. 5 (1943) pp. 115–33.

30. Indeed, Simon received the 1978 Nobel Prize in Economics, and his 1947 book, *Administrative behaviour* (Macmillan, New York), was explicitly mentioned in the prize announcement.

31. As noted in McCorduck, *Machines who think*, p. 109ff.

32. J.M. Brady, 'Report on a visit to the U.S.', Essex University, 1975 (edited version of a report submitted to the Science Research Council, Sept. 1975), outlines the funding situation for Artificial Intelligence in the United States. McCorduck, *Machines who think*, p. 117ff., discusses the 'no strings attached' role of the Air Force funding in allowing the work of Simon and Newell to get started.

33. For attacks on the field see: M. Taube, *Computers and common sense: the myth of thinking machines* (Columbia University Press, New York, 1961); H.L. Dreyfus, *What computers can't do: the limits of Artificial Intelligence* (Harper and Row, New York, 1972); and J. Weizenbaum, *Computer power and human reason* (Freeman, San Francisco, 1976).

34. N. Elias, 'Scientific establishments' (fn. 12).

35. P. Armer, 'Attitudes toward intelligence machines' in E.A. Feigenbaum and J. Feldman (eds), *Computers and thought* (McGraw Hill, New York, 1963), pp. 389–405. This was also commented upon in interviews: R.A. Brooker, Essex University, 28/3/79; and A.M. Uttley, Sussex University, 21/3/79.

36. A.M. Turing, 'Intelligent machinery', (1947) in B. Meltzer and D. Michie (eds), *Machine intelligence 5* (Edinburgh University Press, Edinburgh, 1969), pp. 3–23; and A.M. Turing, 'Computing machinery and intelligence', *Mind*, vol. 59 (1950), pp. 433–60.

37. R.J.W. Craik, *The nature of explanation* (Cambridge University Press, Cambridge, 1952), p. 57.

38. W.R. Ashby, *Design for a brain* (Wiley, New York, 1952); and

An introduction to cybernetics (Methuen, London, 1965).

39. W.G. Walter, 'An imitation of life', *Scientific American*, vol. 182, no. 5 (1950), pp. 42–5; 'A machine that learns', *Scientific American*, vol. 185, no. 2 (1951), pp. 60–3.

40. F.H. George, *The brain as a computer* (Pergamon Press, Oxford, 1961).

41. The RATIO club was discussed in an interview: A.M. Uttley, Sussex University, 21/3/79; and is also discussed by McCorduck, *Machines who think*, p. 59.

42. The proceedings of this conference are published in *Mechanisation of thought processes* (HMSO, London, 1960).

43. J. McCarthy, 'Programs with common sense', ibid., pp. 75–84.

44. D. Michie, *On machine intelligence* (Edinburgh University Press, Edinburgh, 1974), pp. 37, 51, and 66–7; where he states: 'It was from my personal association with Turing during the war and early post-war years that I acquired my interest in the possibilities of using digital computers to simulate some of the higher mental functions that we call "thinking".'

45. This, and much of the following information pertaining to Michie's activities is derived from interviews with Michie (Edinburgh University, 29 and 31/8/78), backed up by other available sources.

46. For example, D. Michie, 'The effect of computers on the character of science', *University of Edinburgh Gazette*, vol. 34 (Oct. 1962), pp. 23–8; and 'The computer revoluton: where Britain lags behind', *The Scotsman*, 12/7/63.

47. *Computing science in 1964*, A Pilot Study of the State of University Based Research in the U.K., prepared for the Research Grants Committee by Dr. Donald Michie (The Science Research Council, London, 1965).

48. Interview: Lord Halsbury, London, 18/12/79.

49. Council for Scientific Policy/University Grants Committee, *A Report of a Joint Working Group on Computers for Research* (HMSO, Cmnd. 2883, London, 1966).

50. Interview: Lord Halsbury, London, 18/12/79.

51. Reconstructed from information in the archives of the Department of Artificial Intelligence, and the Machine Intelligence Research Unit, Edinburgh University, and from interviews.

52. J. McCarthy, 'Review of the Lighthill report', *Artificial Intelligence*, vol. 5 (1974), pp. 317–22.

53. R.M. Burstall and R.J. Popplestone, 'POP-2 Reference Manual' in E. Dale and D. Michie (eds), *Machine Intelligence*, vol. 2 (Oliver and Boyd, Edinburgh, 1968), pp. 207–46.

54. Interviews: B. Meltzer, Edinburgh University, 1/8/77 and 23/1/79.

55. N.L. Collins and D. Michie (eds), *Machine Intelligence*, vol. 1 (Oliver and Boyd, Edinburgh, 1967). There have been nine workshops in the series to date.

56. E.W. Elcock, 'Report of the SRC Computer Research Group', Aberdeen University, May 1970; and interviews: A.M. Murray, Aberdeen University, 9/11/78; P.M.D. Gray, Aberdeen University, 23/7/79; and J.M. Foster, R.R.E. Malvern, 19/6/79.

57. Information on the activities of Clowes derives from a letter: M.B. Clowes, 23/11/78; and interview: M.B. Clowes, Sussex University, 22/3/79.

58. Interview: B. Meltzer, Edinburgh University, 23/1/79.

59. Interview: P.J. Hayes, Essex University, 27/3/79.

60. These conferences, 'The International Joint Conferences on Artificial Intelligence', have been one of the main organs of communication in the area, and a major outlet for publications in the area.

61. R.L. Gregory, *Eye and brain* (Weidenfeld and Nicholson, London, 1966).

62. Information on the robot is derived from numerous documentary sources in the archives of the Department of Artificial Intelligence, and the Machine Intelligence Research Unit, and from various interviews.

63. Interview: J.A.M. Howe, Edinburgh University, 24/1/79.

64. As was evident from various interviews, including one with S. Michaelson, the director of the Computer Science Department at Edinburgh University, on 18/7/79.

65. Lord Swann commented on the new departures in an interview: London, 10/10/79. At a press conference in connection with an open day for industrialists, he is reported as commenting that Edinburgh was among the top two centres of research, despite having no costly 'big science' projects. *University of Edinburgh Bulletin*, vol. 8, no. 2 (Oct. 1971).

66. Interview: R.L. Gregory, Bristol University, 25/7/79.

67. Conversational Software Ltd., launched in 1970.

68. Interview: D. Michie, Edinburgh University, 31/8/78. He commented that at one stage some 20 sources were involved.

69. *Computing Science Review* (Science Research Council, London, 1972), p. 17.

70. Ibid., p. 19. Compare also D. Michie, 'Schools of thought about AI', *University of Edinburgh, School of Artificial Intelligence, Experimental Programming Report*, no. 32 (1973).

71. *Artificial Intelligence: a paper symposium* (The Science Research Council, London, 1973), p. i.

72. This review was carried out by a Special Committee, chaired by Prof. N. Feather, and set up by the University Court. It reported in 1973, and extracts were made available to me by C.H. Stewart, at that time Secretary to the University, in a letter, 18/7/78.

73. As was evident from interviews and also from the Departmental Newsletter.

74. J. McCarthy and P.J. Hayes, 'Some philosophical problems from the standpoint of Artificial Intelligence' in B. Meltzer and D. Michie (eds), *Machine Intelligence*, vol. 4 (Edinburgh University Press, Edinburgh, 1969).

75. G.D. Plotkin, 'Automatic methods of inductive inference', unpub. PhD diss., University of Edinburgh, 1971.

76. Interview: R.A. Kowalski, Imperial College, London, 15/6/79.

77. From the Departmental Newsletter.

78. From the Minutes of the Round Table, a sort of departmental board.

79. Interview: Lord Swann, London, 10/10/79.
80. Various interviews.
81. D. Michie, 'A six year project to develop an intelligent problem solving system' (modified version of a seven-year project proposal to the SRC, Edinburgh University, 1972).
82. P.A. Ambler *et al.*, 'A versatile computer controlled assembly system', in *Proc. Third International Joint Conference on Artificial Intelligence* (Stanford, 1973), pp. 298–307; and *Artificial Intelligence*, vol. 6 (1975), pp. 129–56.
83. Sir James Lighthill, 'Artificial Intelligence: a general survey' in *Artificial Intelligence* (SRC, London, 1973), p. 7.
84. As is evident in the *AISB Newsletter*, *SIGART* (The American Equivalent), and in numerous other places.
85. In the public press there were articles in *The Times Higher Education Supplement*, 1/6/73 and 8/6/73, in *Science*, vol. 180 (June 1973), pp. 1352–3, and in the *New Scientist*, 22/2/73 and 1/3/73. There was a public debate on 4 July, 1973, at the Royal Institution, Albemarle Street, London, under the chairmanship of Sir George Porter, between Sir James Lighthill and Professors R.L. Gregory, J. McCarthy, and D. Michie on Sir James's theme: 'The general purpose robot is a mirage'. This was televised by the BBC as one of its 'Controversy' series and broadcast on 30 August 1973. Later, videorecordings of the debate were taken to the United States as the 'Lighthill Tapes', where they were shown to packed audiences around the centres of Artificial Intelligence.
86. Lighthill, 'Artificial Intelligence: a general survey', p. 1 (fn. 83).
87. In *Proposed new initiatives in computing and computer applications* (Science Research Council, Swindon, March 1979), p. 8, there is the comment: 'The Panel has no doubt that the reluctance of the present community to take up the challenge (of industrial robots research) is due at least in part to the general discouragement of Artificial Intelligence which took place in this country several years ago and that it is now up to SRC to take steps to remedy the situation.'
88. Reported in *AISB*, vol. 20 (July 1975), and in *The Times Higher Education Supplement*, 14/3/75, p. 9.
89. It is impossible to give definite figures in brief for Artificial Intelligence funding because of the multiplicity of sources, and because of the difficulty in deciding what money went to specifically Artificial Intelligence research of the type discussed in this paper; the SRC, for instance have sometimes included under the category of machine intelligence work that is clearly along cybernetic lines.
90. Sources detailed in fn. 56.
91. Interview: R.A. Brooker, Essex University, 27/3/79.
92. Interview: D. Michie, Edinburgh University, 31/8/78.
93. D. Michie, 'Machine intelligence in the cycle shed', *New Scientist*, vol. 57 (22 Feb. 1973), pp. 422–3.
94. Interview: Lord Swann, London, 10/10/79.
95. E.W. Elcock and D. Michie (eds), *Machine Intelligence*, vol. 8, (Wiley, New York, 1977).
96. J.E. Hayes, D. Michie and L.I. Mikulich (eds), *Machine Intelligence*, vol. 9 (Ellis Horwood, Chichester, 1979).

97. Michie organised the 'AISB Summer School on Expert Systems', held at Edinburgh University, July 1979.
98. A. Bundy *et al.*, *Artificial Intelligence: an introductory course* (Edinburgh University Press, Edinburgh, 1978).
99. T. Winograd, *Understanding natural language*, (Edinburgh University Press, Edinburgh, 1972).
100. Various interviews, and fn. 27.
101. J.A. Robinson, 'A machine oriented logic based on the resolution principle', *Journal ACM*, vol. 12, (1965), pp. 23–41.
102. Fleck, 'The structure and development of Artificial Intelligence', pp. 44–8 (fn. 3).
103. D. Warren, 'PROLOG on the DEC system-10', paper presented at the AISB Summer School on Expert Systems, Edinburgh University, July 1979, notes some applications. The position of theorem proving was also discussed in several interviews.
104. Interview: R.M. Burstall, Edinburgh University, 10/8/77.
105. Interview: R.A. Brooker, Essex University, 27/3/79.
106. Interviews: J.M. Brady, Essex University, 26 & 28/3/79; P.J. Hayes, Essex University, 27/3/79.
107. J.E. Doran, 'Knowledge representation for archaeological inference', in E.W. Elcock and D. Michie (eds); *Machine intelligence*, vol. 8 (Wiley, New York, 1977), pp. 433–54.
108. Interviews: B. Anderson, Essex University, 30/7/79; Y. Wilks, Essex University, 8/8/79.
109. Information on movements through Essex was derived from interviews with Brady (see fn. 106), backed up by other sources.
110. Information on developments at Sussex University was derived from various interviews and documentary sources, in particular the interview: N.S. Sutherland, Sussex University, 21/9/79.
111. For example, A.M. Uttley, 'Simulation studies of learning in an informon network', *Brain Research*, vol. 102 (1976), pp. 37–53.
112. M.B. Clowes, 'On seeing things', *Artificial Intelligence*, vol. 2 (1971), pp. 79–116.
113. This organisation is frequently commented on in the *Annual Reports*, Sussex University, and is discussed by the first Vice Chancellor, Asa Briggs, in his 'Drawing a new map of learning' in D. Daiches (ed.), *The idea of a new university: an experiment in Sussex* (Deutsch, London, 1964), pp. 60–80.
114. Letter: Lord Briggs, 28/11/79.
115. 'Working party on School of Cognitive Studies', University of Sussex, June 1970.
116. Interview: M.A. Boden, Sussex University, 19/9/79.
117. M.A. Boden, *Artificial intelligence and natural man* (Harvester Press, Hassocks, 1977).
118. Interview: A. Sloman, Sussex University, 20/3/79.
119. A. Sloman, *The computer revolution in philosophy* (Harvester Press, Hassocks, 1978).
120. Interview: P.N. Johnson-Laird, Sussex University, 20/9/79.
121. For instance, approximately one-quarter of the MSc computer

science courses outlined in *Graduate Studies 1974–75* (CRAC, Cambridge, 1974), mentioned Artificial Intelligence, and it is often mentioned as an acceptable interest in job advertisements.

122. J.M. Brady, 'A glimpse of the future of AI', text of a lecture delivered at 'Computing 79', Sydney, Australia, August 1979, p. 2.

123. B.C. Griffith and N.C. Mullins, 'Coherent social groups in scientific change', *Science*, vol. 177 (1972), pp. 959–64.

124. I do not want to argue a 'great man' view of history here. The situation can be interpreted in terms of the emergence of socially defined possibilities which in the event were exploited by Michie. If Michie had not been, development in Artificial Intelligence in Britain would still have taken place with other people stepping in to a greater extent. Perhaps what is required is a 'great opportunities' view of history, with the emphasis on the socially given possibilities rather than on the people who exploit them.

125. T. Shinn, 'Scientific disciplines and organisational specificity: the social and cognitive configuration of laboratory activities', in N. Elias, H. Martins and R. Whitley (eds), *Scientific establishments and hierarchies*, Sociology of the Sciences, vol. VI (Reidel, Dordrecht, 1982), pp. 239–64.

126. See for instance, P. Kraft, *Programmers and managers: the routinisation of computer programming in the United States* (Springer Verlag, New York, 1977).

127. M.J. Mulkay, in *The social process of innovation* (MacMillan, London, 1972), concludes: '. . . intellectual migration, whether into established networks or into virgin areas, will normally be led by mature researchers of known repute. These men use their eminence to attract funds and graduate students; and in various ways they try to use their existing knowledge and techniques as a point of departure for the construction of the new intellectual framework' (p. 54). The pattern of development of Artificial Intelligence in Britain certainly conforms with the first part of what Mulkay writes, but it is difficult to square it with the second part. It would appear that the extent of articulation of the paradigmatic structure in the United States preempted attempts by migrants in Britain to use their existing knowledge in the construction of new intellectual frameworks, and left them scope only for negotiation over the organisational aspects.

128. See N.S. Sutherland, 'Neuroscience versus cognitive science', *Trends in Neurosciences*, vol. 2, no. 8 (1979), pp. i–ii, which discusses some of the particulars of this dominance.

129. Lighthill, 'Artificial Intelligence: a general survey', p. 1 (fn. 83).

130. For example, Sir Geoffrey Allen (Chairman of the Science Research Council), in a talk in the Department of Liberal Studies in Science, the University of Manchester, 17/5/79.

131. C.E. Shannon, 'Automatic chess player', *Scientific American*, vol. 182, no. 2 (1950), pp. 48–51.

132. See fn. 33.

133. Turing, 'Intelligent machinery' (fn. 36).

134. Fleck, 'The structure and development of Artificial Intelligence', pp. 138–9 (fn. 3).

135. Such an impression was given by those in Artificial Intelligence (for example, interview: A. Bundy, Edinburgh University, 17/7/79) and those outside, in the course of interviews.

136. Dreyfus, *What computers can't do*, is the classic example here (fn. 33).

137. Weizenbaum, *Computer power and human reason* (fn. 33).

138. Taube, *Computers and common sense* (fn. 33).

139. E.W. Dijkstra, 'Programming: from craft to scientific discipline', *Proc. International Computing Symposium, 1977* (Liège, Belgium, April 1977), pp. 23–30.

140. The Japanese success in catching up with and then overtaking the leading Western economies in shipbuilding, automobiles, consumer electronics, capital goods, and general manufacturing hardly needs further comment here, but it certainly added substance to the threat posed by the 5th Generation Project.

141. *Summer 1980 JIPDEC Report*, (Japan Information Processing Development Center, Tokyo, September 1980).

142. D'Agapeyeff was a founder of CAP Ltd., London, at one time one of the biggest software houses in Europe.

143. During a talk entitled 'Challenges for expert systems', given at the Expert Systems: 81 Conference, London, December 1981.

144. For instance: A. d'Agapeyeff, 'Report to the Alvey Directorate on a short survey of expert systems in UK business', *Alvey News*, Supplement to Issue No. 4 (April 1984).

145. See Figure 3.1 in section 3.5 of this chapter.

146. McCorduck wrote a history of AI in the US: *Machines who think* (fn. 18).

147. E.A. Feigenbaum and P. McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's challenge to the world* (Michael Joseph, London, 1983), quote taken from the prolog.

148. Department of Industry, *A programme of advanced information technology*, The report to the Alvey Committee (HMSO, London, 1982).

149. Quoted in J. Alvey, 'UK response to the Fifth Generation', *Electronics and Power*, May (1983), pp. 387–9.

150. See: 'The Fifth Generation: taking stock', *Science*, 30 Nov. (1984), 1061–3.

151. Information derived from research in progress on an ESRC funded project: 'The effective management of available expertise in Artificial Intelligence', held by D. Edge and J. Fleck (database at Jan. 1987).

152. From the same source.

153. For instance, see the survey by d'Agapeyeff, 'Report to the Alvey Directorate' (fn. 144).

154. K. Cornwall-Jones, PhD research in progress, Science Policy Research Unit, University of Sussex, 1986.

155. For instance, A. Bundy *et al.* (eds), *A catalogue of Artificial Intelligence tools* (Science and Engineering Research Council, Swindon, Spring 1983).

156. As can be seen from the papers presented in the annual series of conferences 'Expert Systems: 81–86'.
157. See *Summer 1980 JIPDEC Report*, p. 15 (fn. 141).
158. M. de Mey, *The cognitive paradigm* (Reidel, Dordrecht, 1982).
159. Compare Descartes's use of the hydraulics metaphor.
160. For example, I. Aleksander, whose work on a pattern recognition chip was based more on a neural net tradition than on the list processing AI approach, has been successful in claiming this as bona fide AI research.
161. For a discussion of the various positions taken by protagonists in these debates, see: J. Fleck, 'Artificial Intelligence and industrial robots: an automatic end for utopian thought?' in E. Mendelsohn and H. Nowotny (eds), *Nineteen eighty-four: science between utopia and dystopia* (Reidel, Dordrecht, 1984), pp. 189–231.
162. For instance Searle certainly created a stir with his 1984 Reith Lecture series and his 'chinese box' critique of AI: J. Searle, *Minds, brains and science* (British Broadcasting Corporation, London, 1984).
163. See B.J. Rooney, 'A Survey of the Technical, Social, and Possible Impact Areas of Expert systems Technology', unpub. MSc diss., University of Aston, 1983.
164. Cornwall-Jones, research in progrss (fn. 154).
165. Seminar presentation by K. Guy, of Science Policy Research Unit, at Edinburgh University, 11 March 1986.
166. This educational effort is amply illustrated by the many training seminars, tutorials, and advanced courses on offer from the various bodies in the area.