

‘What is science for? The Lighthill report and the purpose of artificial intelligence research’

Jon Agar

STS, UCL

Abstract. This paper uses a case study of a 1970s controversy in artificial intelligence (AI) research to explore how scientists understand the relationships between research and practical applications, part of a project that seeks to map such relationships in order to enable better policy recommendations to be grounded empirically through historical evidence. In 1972 the mathematician James Lighthill submitted a report, published in 1973, on the state of artificial intelligence research underway in the United Kingdom. The criticisms made in the report have been held to be a major cause behind the dramatic slowing down (subsequently called an ‘AI winter’) of such research. This paper has two aims, one narrow and one broad. The narrow aim is to inquire into the causes, motivations and content of the Lighthill report. I argue that behind James Lighthill’s criticisms of a central part of artificial intelligence was a principle he held throughout his career that the best research was tightly coupled to practical problem solving. I also show that the Science Research Council provided a preliminary steer to the direction of this apparently independent report. The broader aim of the paper is map some of the ways that scientists (and in Lighthill’s case, a mathematician) have articulated and justified relationships between research and practical, real world problems, an issue previously identified as central to historical analysis of modern science. The paper therefore offers some deepened historical case studies of the processes identified in Agar’s ‘working worlds’ model.

Most science in the twentieth century was oriented, in some way, towards what I have called “working worlds”, ‘arenas of human projects that generate problems’, of which important examples include the preparation, mobilisation and maintenance of military systems and the maintenance of the human body in sickness and health.¹ Working world problems are typically not solvable directly, nor are they immediately tractable by the scientific methods. Science therefore builds ‘manageable, manipulable, abstracted’ representatives, such as models, data sets, experimental proxies, that in their simplified form stand in for the world.² Working worlds are a cause of, and a bind and opportunity for, science. They are an opportunity because problem-solving provides a justification (for science) for resources and the satisfaction (for scientists) of effectively intervening in the world, amongst other reasons. They are a bind because problem articulation can be incessant and overwhelming. Indeed, the social institutions and norms of science have been shaped to provide scientists with a measure of autonomy and protection.

In *Science in the Twentieth Century and Beyond*, working worlds were very much an analyst’s category. I found them a useful label for a pervasive feature of the landscape within which science necessarily was placed. Since publication, in 2012, I have been thinking further about the concept, not least in response to constructive criticism. I am intending to strengthen and refine the analytical framework by answering such questions as: how many working worlds are there? How far back can the analysis be pushed back usefully? Are there working worlds that do not generate science? Are there insulated areas of science with no connection to working worlds? Are there science policy consequences, and how might they be implemented and tested? Answering such questions would be a path towards sharpening the analytical term through trials against empirical evidence.

But there is a second way forward, one which explores actors’ categories, and one which I want to take a step towards in this paper. Scientists, as historical actors, have perceived, interpreted,

¹ Jon Agar, *Science in the Twentieth Century and Beyond*, Cambridge: Polity, 2012, p. 3.

² Agar, *op. cit.*, p. 4.

represented and intervened into working worlds. How they articulate these relationships can be the focus of historical study. The second project therefore is to map, historically, these kinds of relationships as they were understood by scientists. If ‘working worlds’ is the analyst’s category, then these relationships are the actors’ categories, and mapping their variety is an important step in understanding how science has in the past related to the worlds’ problems and how in the future they might relate better.

My case study concerns the British applied mathematician Sir James Lighthill. While he was highly productive in the field of fluid dynamics, he is perhaps best known more widely as the author of a critical report on the state of artificial intelligence in the early 1970s.³ Indeed, in the following I will demonstrate that our historical understanding of the motivation and content of the Lighthill report needs to be revised in the light of fresh archival discoveries. Specifically, I show that Lighthill received a steer from the commissioning research council, but also that Lighthill’s analysis of the field of artificial intelligence was guided by a broader and deeper commitment he held towards the relationships science (and mathematics) should have towards problem solving. My further interest in the Lighthill report episode, therefore, is as a case study that reveals how scientists – in this case the mathematician Lighthill and his opponents among the community of artificial intelligence science - have had to articulate and justify a relationship between their work and the world’s problems.

One way of summarising my case study is to say, first, that Lighthill embraced a “pro-working worlds” view of science, and explicitly argued that good science (not least his own applied mathematics) not only ultimately addressed problems but was fundamentally oriented towards them. However, he must be placed on an extreme of a spectrum of attitudes, certainly within mathematics (we might put G.H. Hardy, who proudly defended the seriousness and uselessness of pure mathematics⁴, on the other wing) and also within science. Second, therefore, we might look for

³ Published as part of Science Research Council, *Artificial Intelligence: a Paper Symposium*, 1973.

⁴ ‘I have never done anything “useful”. No discovery of mine has made, or is likely to make, directly or indirectly, for good or ill, the least difference to the amenity of the world. I have helped to train other

and study “anti-working worlds” articulations of science. These articulations are varied and subtle, and an important goal of the broader project is to map them. Just as we follow controversies because science’s assumptions are brought to the surface, so we historians and sociologists can study scientists’ pro and anti-working world positions, as they provoke each other, to understand them better. In this paper I explore how Lighthill’s intervention prompted scientists, in particular, the Edinburgh computer scientist Donald Michie, to reflect on to what extent his science did or did not, should or should not, respond to practical problems.

The field of Artificial Intelligence has followed a pattern of boom and bust ever since its establishment in the 1950s. In his history of AI, Daniel Crevier notes sharp American funding cut backs in 1974, and then turns to summarise what had happened in Britain:

American researchers were not the only ones to feel the ebbing of the tide. A scathing report to their government on the state of AI research in England [sic] devastated their British colleagues. Its author, Sir James Lighthill, had distinguished himself in fluid dynamics and had occupied Cambridge University’s Lucasian Chair of Applied Mathematics, presently held by Stephen Hawking ... Sir Lighthill [sic] pointed out to the Science Research Council of the UK that AI’s “grandiose objectives” remained largely unmet. As a result, his 1973 report called for a virtual halt to all AI research in Britain.⁵

This is the familiar account.⁶ AI is charged with hubris, grandiosity, claiming too much, reaching for the stars. Brought down to earth with a bump it had to slowly rebuild. The cycle gives us what were

mathematicians, but mathematicians of the same kind as myself, and their work has been, so far at any rate as I have helped them to it, as useless as my own. Judged by all practical standards, the value of my mathematical life is nil’. G.H. Hardy, *A Mathematician’s Apology*, Cambridge: Cambridge University Press, 1940, p. 49.

⁵ Daniel Crevier, *AI: the Tumultuous History of the Search for Artificial Intelligence*, New York: Basic Books, 1993, p. 117

⁶ See, also, for example, Nils J. Nilsson, *The Quest for Artificial Intelligence: a History of Ideas and Achievements*, Cambridge: Cambridge University Press, 2010, p. 204, p. 282. There is a surprising absence in the secondary literature of survey histories of artificial intelligence written by professional historians of science. There is, however, a growing corpus of excellent, more specific studies, for example: Alison Adam, *Artificial Knowing: Gender and the Thinking Machine*, London: Routledge, 1998, Paul N. Edwards, *The Closed World: Computers and the Politics of Discourse in Cold War America*, Cambridge, MA: MIT Press, 1996.

subsequently called 'AI winters', during which funding, and therefore research, froze. It is an internal, autochthonous historiography, part of the received culture of AI. It is both self-critical (we promised too much) and self-bolstering (but look at how ambitious our aims are). It blames both ancestors and outsiders. It is sometimes done with humour. The problem, said Douglas Lenat, the leader of the brute-force AI Cyc project, was the jerks in funding. By which he meant

... the inevitable up-and-down nature of the funding streams themselves. As a former physics student, I remember that the first, second, and third derivatives of position with respect to time are termed velocity, acceleration, and jerk. Not only are there changes in funding of AI research and development (velocity), and changes in the rate of change of funding (acceleration), there are often big changes in the rate of those accelerations year by year—hence what I meant by jerks in funding.⁷

The funder, in Britain, in the 1970s was the Science Research Council (SRC). The cause of the first AI winter was a critical report commissioned by the SRC, written by the mathematician James Lighthill. Specifically, Crevier argues that the cause of the downturn was Lighthill's criticism of AI's grandiose objectives. One of Lighthill's obituarists extended the list of complaints, writing that

...responding to a Science Research Council invitation, he infuriated the computing lobby by advising against investment in artificial intelligence because he believed the concept difficult, costly, and probably ill-founded. Many now believe him to have been right.⁸

Stephanie Dick, 'AfterMath: the work of proof in the age of human-machine collaboration', *Isis* (2011) 102, pp. 494-505. Nathan Ensmenger, 'Is chess the drosophila of artificial intelligence? A social history of an algorithm', *Social Studies of Science* (2012) 42, pp. 5-30.

⁷ D.B. Lenat, 'The voice of the turtle: Whatever happened to AI?', *AI Magazine* (2008) 29(2), pp. 11-22, p. 16.

⁸ Anthony Tucker, 'Sir James Lighthill. Maths in motion', *The Guardian* 21 July 1998. The obituary is riddled with errors, not least consistently calling Lighthill 'Lightman'.

While Lighthill's report did indeed cause AI research to stall – the criticisms resonated beyond the UK's shores - I suggest that Lighthill's critique, or rather the principled motivation of it, has been misunderstood. In the following I introduce James Lighthill, summarise his report and the responses of his critics, focussing in particular on Donald Michie, the Edinburgh machine intelligence pioneer most directly affected by the report. I will demonstrate that Lighthill made a career-long commitment to the idea that the best science (and mathematics) was that produced from a tight relationship between research and practical problem-solving. This commitment can be found before and after the report, but crucially also underpins Lighthill's thinking about artificial intelligence in the report itself. Finally, I step back and discuss the relevance of this episode to the broader project: to understand the contested history of science as profoundly shaped by an orientation towards practical problem-solving, deploying and developing the "working worlds" historiography. Lighthill and Michie, as I will show, had different views on this issue. This paper therefore has two aims, a narrow one to uncover why Lighthill offered such a critical report on artificial intelligence, and a broader one to understand the ways that scientists have viewed the relationship of scientific research to real world problem solving in the past.

Who was Lighthill?

Michael James Lighthill was born in 1924, won a scholarship to Trinity College, Cambridge, at age 15 (alongside the other prodigy Freeman Dyson). During the Second World War he worked at the National Physical Laboratory on supersonic and hypersonic aerodynamics. Afterwards, he taught mathematics at Manchester University, where he 'ran one of the most powerful and inventive fluid dynamics groups ever formed anywhere'⁹ and researched jet engines, including the suppression of noise from the prototype Boeing 707.¹⁰ In 1959 he was appointed director of the Royal Aircraft

⁹ David George Crighton, 'Sir James Lighthill FRS', *The Independent*, 22 July 1998.

¹⁰ His 1952 paper launched the field of 'aeroacoustics' ('sound generated aerodynamically') and, says Crighton, remarkable for two reasons: first, Lighthill delayed publication for sixteen months so that he could rewrite in a

Establishment (RAE), Farnborough. His work on supersonic wing shape was used in the design of Concorde. He also encouraged RAE and the Post Office to develop proposals for communications satellites.¹¹ In 1964 he returned to academia as a Royal Society professor at Imperial College, before moving to Cambridge as Lucasian Professor of Mathematics in 1969, succeeding Paul Dirac. From 1979 he was Provost of University College London, where he remained for a decade. At UCL he ‘stressed the need for universities to produce graduates in useful subjects such as building’. In 1998, Lighthill he died while swimming around the island of Sark, a feat he had completed four times previously.¹²

In terms of his mathematics, Lighthill specialised in fluid dynamics, but in an original and distinctive way. Some of his papers – notably the 1952 paper on aeroacoustics and the 1956 long paper, written in honour of G.I. Taylor, on nonlinear acoustics – were progenitors of new fields. Each started in one practical context, but found application in a wide variety of other contexts: the 1952 paper was on jet engine noise but has been applied to understanding the sun’s corona, while the applications of the 1956 paper include ‘kidney-stone-crushing lithotripsy machines, ... flood waves in rivers and traffic flow on highways’.¹³ Conversely, his mathematics could also spring from analyses of diverse situations: his biofluidynamics, which spanned blood flow, breathing, bird and insect flight, and the swimming of fish, is a good example.

Before I turn to Lighthill’s report on AI, I would like to emphasise some aspects of his career that will be important. First, Lighthill was an advocate for applied mathematics. Not only did his work start and end with applications, but he promoted the approach institutionally too. In 1965, for example,

form understandable to aeroengine designers and, second, the ‘paper neither contains nor needed as much as a single reference to any prior work’

¹¹ MJL D4, untitled document beginning ‘Ladies and Gentleman’ is his resignation speech to RAE staff, and includes a review of achievements and developments.

¹² Another obituary says 6. Also from this *Telegraph* obituary: he was first person to swim around Sark, in 1973. He used a “two-arm, two-leg backstroke, thrusting with the arms and legs alternately. ‘Sir James Lighthill, *Daily Telegraph*, 20 July 1998.

¹³ Crighton, op. cit.

he established a new Institute of Mathematics and its Applications, borne of frustration with the pure emphasis of the London Mathematical Society.

Second, in 1966 Lighthill had been invited on to a working party to review the UK fusion research programme at Culham.¹⁴ Lighthill, along with the electronics expert F.E. Jones¹⁵, viewed the working party's report as 'half-hearted', broke ranks and offered a much more critical minority report, one which represented better the 'views of informed industrialists and scientists outside the Authority [the United Kingdom Atomic Energy Authority]'.¹⁶ Lighthill and Jones argued that no work on the production of power through fusion reactors was justifiable, and they rejected the arguments presented in the main report about the 'likely importance [of fusion power research and development] to science and technology in general' and the need for 'national and international effort on plasma physics and fusion research'. This willingness to strike boldly and independently— to attempt to kill fusion research in this case - was surely noted within government circles. It is also perhaps significant that Lighthill, with Jones, were not only willing to end UK fusion research, but they also suggested that Culham's 120 scientific and 130 other staff might be redirected to problem-solving work on 'efforts to raise the level of efficiency of British industry', such as by conducting

¹⁴ MJL K686. Vick to Lighthill, 25 January 1966.

¹⁵ Francis Edgar Jones – Frank Jones – had worked during the war at the Telecommunications Research Establishment (TRE), leading the Oboe development team, his managerial competence providing a pragmatic complement to the more disorganized but inventive A.H. Reeves. Unlike many of the TRE high-fliers, he did not return to academic research and instead supported the attempts to build a post-war role for TRE. In 1953, six years before Lighthill himself arrived, Jones joined the Royal Aircraft Establishment, becoming Deputy Director (Equipment). In 1956 he was recruited by the electronics company Mullard Ltd where he remained until retirement in 1972. Like Lighthill he championed the application of science. For example, with Mullard's John Bunton he persuaded the Royal Society to mint a new prize in 1967, the Mullard Medal, because both 'felt that a medal for technology and applied science on the same level as the Copley Medal could make a contribution in influencing scientific effort in that direction'. He was an accomplished builder of networks between industry, government and science. While his Royal Society obituarist does not mention Jones' role in the Culham report, he does note Jones' chairing of the more publicised output of the Working Group on Migration of Scientists, the 1967 'Brain Drain' report. George G. MacFarlane and Cyril Hilsum, 'Francis Edgar Jones, 16 January 1914 - 10 April 1988', *Biographical Memoirs of Fellows of the Royal Society* (1988) 35, pp. 179-199, p. 194, p. 196.

¹⁶ MJL K694. 'Culham Laboratory Review Panel. Minority report by Dr F.E. Jones and Professor M.J. Lighthill', undated (1966). They argued that no work on fusion power was justifiable.

‘intensive surveys of the technological problems of particular firms or groups of firms and work out comprehensive plans for improvements’.¹⁷

The Lighthill review

In September 1971, the chair of the Science Research Council, Brian Flowers, contacted Lighthill.

‘There are few subjects which at any particular time strike one as having a very special potential for being pervasively important’, he wrote, but ‘Artificial Intelligence seems to me to be such a field, overlapping as it does with neurobiology, psychology, linguistics, and computer-aided learning, not to mention mathematics and computer science proper’. Flowers then got to the critical point:

These subjects, and artificial intelligence itself, are highly complex, very much the preserve of experts and perhaps sometimes of plausible charlatans. There is a strong tendency for the experts to pursue their particular enthusiasms energetically but narrowly and to ignore, or repudiate any other approach ... It is getting increasingly difficult for the SRC to control this mix of activities and to make properly informed judgments...¹⁸

With this remarkable steer, Flowers asked Lighthill ‘to make a considered appreciation of the subject of artificial intelligence, its achievements, its practitioners, its promise and its needs’. While Lighthill was initially reluctant, by late 1971 and early 1972 he was very busy investigating the state of AI in Britain. He wrote letters to some of those active in AI in the UK, specifically Donald Michie and R.M. Burstall at the Department of Machine Intelligence and Perception, University of Edinburgh, B. Meltzer, also at Edinburgh but in mathematics, M. Clowes and N.S. Sutherland at Experimental Psychology, University of Sussex, E.W. Elcock of Computing Science at University of Aberdeen and J. Foster at Computing Science at the University of Essex. These academics, at four universities, were

¹⁷ MJL K694, ‘Minority report’, p. 7.

¹⁸ Lighthill papers, UCL Collections (hereafter MJL) K668, Flowers to Lighthill, 20 September 1971.

asked their opinions about the potential significance of the field, the value of current research, the adequacy of funding and organisation, and how the field should be developed, as well as an account of their own contribution to the field and current research objectives.¹⁹ A second, wider group of mostly academics in related research areas were consulted about their external views of AI.²⁰

Not all of the responses to Lighthill have been found in the archives so far, although we have more of his replies.²¹ Michie sent at least three detailed letters, and his group (as well as Meltzer's staff and Christopher Longuet-Higgins) met Lighthill in Edinburgh, where they stood up to a 'somewhat intensive battery of probing questions'.²² Lighthill also met with Alan H. Cook, professor of geophysics, chair of the steering group for the Edinburgh AI work, who had briefed Lighthill on his view of the department's 'problems'.²³ Up until this point the letters exchanged between Lighthill and Edinburgh had been friendly.

Lighthill submitted his report in March 1972.²⁴ 'Quite frankly', Lighthill wrote in a following letter, 'I am fully aware that when my report becomes widely available I shall be involved in a great deal of

¹⁹ MJL K668. Draft letter from Lighthill, undated.

²⁰ MJL K668. 'Investigation of artificial intelligence research'. The external list included, among computer scientists: R.A. Brooker (Essex), R.M. Needham (Cambridge), D.C. Cooper (Swansea), P.J. Landin (Queen Mary); and among neurobiology, linguistics and other disciplines: R.L. Gregory (Brian Perception Unit, Bristol), H.C. Longuet-Higgins (Machine Intelligence and Perception, Edinburgh), R.C. Oldfield (MRC Speech and Communications Research Unit, Edinburgh), D. Fry and A.J. Fourcin (Phonetics, UCL), D.M. MacKay (Communications, Keele), S. Brenner (MRC Laboratory of Molecular Biology, Cambridge), J.N. Holmes (Joint Speech Research Unit, General Post Office) and J.J. Kulikowski (UMIST). The rough notes of replies shows that several others were also contacted, including Christopher Strachey (Oxford), Tom Kilburn (Manchester) and D.E. Broadbent (Applied Psychology Unit, Cambridge). In the report itself, Lighthill lists the following: J. Annett, H. G. Barrow, S. Brenner, D. E. Broadbent, R. A. Brooker, O. P. Buneman, R. M. Burstall, A. D. B. Clarke, M. B. Clowes, A. H. Cook, D. C. Cooper, J. E. Doran, J. F. Duke, E. W. Elcock, I. J. Good, C. C. Green, R. L. Gregory, P. J. Hayes, A. L. Hodgkin, J. N. Holmes, J. A. M. Howe, D. H. Hubel, S. Isard, H. Kay, T. Kilburn, J. Kulikowski, D. N. L. Levy, H. C. Longuet-Higgins, D. M. MacKay, D. Marr, J. McCarthy, B. Meltzer, D. Michie, M. Minsky, D. Mollison, E. Moore, J. S. Moore, R. M. Needham, N. J. Nilsson, C. Oldfield, J. V. Oldfield, I. Pohl, R. J. Popplestone, B. Raphael, J. A. Robinson, C. Strachey, N. S. Sutherland, M. M. Swann, H. P. F. Swinnerton-Dyer, D. Wilshaw, T. Winograd.

²¹ I discuss some below. In addition extracts of Christopher Strachey's response, 15 February 1972, is in Add MS 88958/1/383.

²² MJL K669. Lighthill to Michie, 16 February 1972.

²³ MJL K669. Lighthill to Cook, 18 February 1972. Lighthill also took the opportunity to congratulate Cook on his imminent move to Cambridge to become Jacksonian Professor. With Cook ending his direct tie with Edinburgh and moving to Lighthill's university, the briefing could have been particularly frank.

²⁴ MJL K670. Lighthill to Flowers, 23 March 1972. A later estimate of expenses shows that Lighthill had worked 39 whole days, for which he charged £12.60 per day, a total expense of £491.40.

controversy'.²⁵ 'The people closely concerned here have now read your report on artificial intelligence [and there is] general agreement that it is a masterly survey of a complex field', replied Flowers, adding 'though naturally enough not everyone is entirely happy about the lessons you draw for the Council'.²⁶ A small coterie of experts was convened to compose an "office view" that would introduce and steer discussion of the report when it was presented to the SRC.²⁷

Lighthill's report was published by the SRC in 1973 as part of *Artificial Intelligence: a Paper Symposium*, alongside a response by Professor N. Stuart Sutherland of the University of Sussex, and three critical replies by Dr Roger M. Needham of the University of Cambridge, Professor Hugh Christopher Longuet-Higgins and Professor Donald Michie, both of Edinburgh.

Lighthill divided the subject of AI into three categories: A, B and C. For reasons that will become clear, it helps to introduce the categories in the order A, C, B. Category A stood for 'Advanced Automation':

The clear objective of this category of work being to replace human beings by machines for specific purposes, which may be industrial or military on the one hand, and mathematical or scientific on the other.²⁸

Examples were mainly from the application of general-purpose digital computers to extend – and be continuous with - present automation in fields such as control engineering, clerical data processing, pattern and speech recognition, component design and manufacture, cryptography, missile guidance, and logical deduction in mathematics. The work linked to the disciplines of computer science and control engineering, especially.

²⁵ MJL K670. Lighthill to Flowers, 27 April 1972.

²⁶ MJL K670. Flowers to Lighthill, 20 April 1972.

²⁷ This group would be chaired by Flowers and include Lighthill, Dr D.E. Broadbent, Dr R.M. Needham, Professor G.S. Brindley (Institute of Psychiatry, London), Professor R.A. Brooker (Essex) and Professor N.S. Sutherland.

²⁸ James Lighthill, 'Artificial intelligence: a general survey' (Lighthill report), in SRC, *Artificial Intelligence: a Paper Symposium*, London: SRC, 1973.

Lighthill's Category C concerned 'Computer-based CNS [Central Nervous System] research'. Examples included neural nets as models of the brain and nervous system, and other ways of modelling neurobiological and psychological activities such as visual pattern recognition, memory, language use, and the acquisition of knowledge and skills. The work linked to the disciplines of neurobiology and psychology.

A and C were, said Lighthill, quite distinct, and if that was all there was they would 'warrant completely separate treatment in respect of research support, departmental organisation, etc'. A and C were also marked by 'perfectly clear' 'aims', 'practical and technological' in the case of A and 'fundamental, biological' in the case of C.

Category B stood for 'Bridge' but also "Building Robots". 'During the same period [in which A and C progressed, both starting from Turing] a further category "B" of researches has been pursued', reported Lighthill, 'a "bridge" category where aims and objectives are much harder to discern but which leans heavily on ideas from both A and C and conversely seeks to influence them'. Only the existence of B creates AI's claim for 'unity and coherence', argued Lighthill, yet there was a 'widespread feeling ... that progress in this bridge category B has been even more disappointing both as regards the work actually done and as regards the establishment of good reasons for doing such work and thus for creating any unified discipline'. Failures of work, especially in category B, to meet the 'inflated predictions' of AI were the root cause of diminished 'confidence in whether the field of AI has any true coherence'.²⁹ Types of work under category B included mimicking the coordination of eye and hand, visual scene analysis, use of natural language, playing games, and "commonsense" problem solving.

Lighthill speculated about the causes of interest in category B, especially building robots: perhaps 'scientists consider themselves duty bound to minister to the public's general penchant for robots by

²⁹ Lighthill did identify failures in categories A and C. For example, mathematical theorem-proving, category A work, 'is particularly an area where hopes have been disappointed through the power of the combinatorial explosion in rapidly cancelling out any advantages from increase in computer power'

building the best they can?', perhaps 'the stimulus to laborious male activity ... is the urge to compensate for lack of female capability of giving birth to children'? But his main point was that category B work had largely failed.³⁰ It was only successful in abstract play situations and problem solving 'when and only when the programming has taken into account a really substantial quantity of human knowledge about the particular problem domain', in other words where there were human-sourced heuristics to guide.

Concluding, Lighthill noted the 'bimodal distribution of achievement', while forecasting over the next 25 years:

within category A or category C, certain research areas making very substantial further progress, coupled in each case with the forging of far stronger links to the *immediate field of application* than to the supposed bridge activity B. Rising confidence about the work's relevance within the associated *field of application* may add prestige and thence strength to such an area of research, while continued failures to make substantial progress towards stated aims within category B may cause progressive loss of prestige, from which a diminution of funding will ultimately follow even where scientific claims are not always subject to full scientific scrutiny. In due course the overriding significance of the links between each research area and its *field of application* will rupture the always fragile unity of the general concept of AI research.³¹

So in category A (at least the engineering side): 'techniques for Advanced Automation can now be expected to move forward fastest *where research workers concentrate upon practical problems, acquiring for the purpose good detailed knowledge of the technological and economic contexts of the problems chosen*'.³² Likewise, in the other side of category A, mathematics, a 'similar outward-

³⁰ Lighthill included a postscript that dealt with Terry Winograd's PhD on the SHRDLU program, for which Lighthill had high praise, while noting that 'One swallow does not make a summer'. It was successful category B work, although the best further developments would come through category C.

³¹ My emphasis.

³² My emphasis.

looking trend is expected' as work moves to the 'utilisation of far more detailed observation of how mathematicians actually prove theorems!'³³ The general point was that 'chances of success in any one area will be greatly improved through close integration of the researches with the field of application'. For category C 'success will again be related to how closely the work is linked to the fundamental associated disciplines of psychology and neurobiology'. We might note here (although Lighthill does not) that psychology and neurobiology were also closely associated with fields of application. Finally category B would continue to falter as categories A and C thrived. The result would be 'the fission of the field of AI research'.

Critical Responses to Lighthill's Report

Lighthill's report was followed in the 'paper symposium' by responses, from Sutherland, Needham, Longuet-Higgins and Michie. Relabelling 'B' as 'Basic research in AI' rather than 'Bridge' or 'Building Robots', Sutherland argued that:

Lighthill's definition of area B is misleading, that some of his arguments against work in this category are unfounded, that the achievements and promise of the work can be seen in a very different light from that which he presents them and that it is hard to see how work in areas A and C can flourish unless there is a central core of work in area B.³⁴

In conclusion, Sutherland argued that

Lighthill's area B so far from being a bridging area is really the central area of progress in AI that work in this area is worth supporting in its own right and that if it is not supported areas

³³ 'The structuring and utilisation of scientific data bases is another area where good results depend on detailed study of the data's special characteristics. The one part of that field with which the present author has been closely involved, as Chairman of the Steering Committee for the Experimental Cartography Unit of NERC since its inception, affords a good example of this: the structuring of geographically located data is found to demand quite specialised techniques, closely related to the cartographic character of the output.'

³⁴ N. Stuart Sutherland, 'Part II Some Comments on the Lighthill report and on Artificial Intelligence', in SRC, *Artificial Intelligence: a Paper Symposium*, op. cit.

A and C will suffer, both through a dearth of the sort of new concepts produced by workers in area B and also through a lack of trained workers in AI, since area B appears to be the most appropriate training ground for workers in all three divisions of the subject.

He suggested the SRC increase, not withdraw, funds from category B, among a number of other means of support. In other words, Sutherland thought that Lighthill had not appreciated the basic value and achievements of category B. He missed Lighthill's underlying emphasis on practical problems.

Roger Needham, on the other hand, basically supported Lighthill and undermined Sutherland's reply.³⁵ Needham reviewed the claims that some category B work had produced worthwhile 'developments in computing technology', in other words that it could be justified by its 'side effects', understood as developments in computer science rather than practical applications, only to find little (but not no) evidence. Indeed, what technical developments had occurred did so despite, not because, of AI, and 'one could argue that the sooner we forget [the 'pernicious label' AI] the better'.

Christopher Longuet-Higgins's strategy was to broadly agree with Lighthill but seek to defend one – his own - particular corner of research: cognitive science, researching computer programmes as modelling the software of the brain.³⁶ This response defended cognitive science (potentially category B, but one he tactically placed in the safer category C) on scientific not practical grounds.³⁷ However

³⁵ Roger Needham, 'Comment by Dr R.M. Needham', in SRC, *Artificial Intelligence: a Paper Symposium*, op. cit. 'Since I basically agree with Lighthill's conclusions there is perhaps less to say than in Sutherland's commentation.'

³⁶ Longuet-Higgins, 'Comment by Professor H.C. Longuet-Higgins FRS', in SRC, *Artificial Intelligence: a Paper Symposium*, op. cit. 'Sir James places in his category C all the artificial intelligence work which he regards as scientifically promising, and refers to this category as Computer based studies of the central nervous system. In so doing he aligns himself with those of us who hold that the main justification for artificial intelligence is the light it can throw upon human intellectual activity. But his chosen heading, and some of his later remarks, indicate that he attaches more significance to work on the hardware of the brain than to work on its software. This is the only point on which I want to take issue with him.'

³⁷ 'In short whatever the technological prospects of artificial intelligence, its principal scientific value, in my view, is that it sets new standards of precision and detail in the formulation of models of cognitive processes, these models being open to direct and immediate test.'

there was an interesting sting in the tail, or at least a warning that without some category B work then category A work could lead to dangerous developments:

Finally, perhaps one should say a word about the main point of disagreement between Lighthill and Sutherland. Professor Sutherland's redefinition and reinstatement - of Lighthill's category B as basic artificial intelligence has my sympathy, because although I hold no particular brief for bridging activities as such, I do think that there is a place in artificial intelligence for studies which are addressed to the general problems which have been found to recur in many different areas of cognitive science. The mathematician's ability to discover a theorem, the formulation of a strategy in master chess, the interpretation of a visual field as a landscape with three cows and a cottage, the feat of hearing what someone says at a cocktail party and the triumph of reading one's aunt's handwriting, all seem to involve the same general skill, namely the ability to integrate in a flash a wide range of knowledge and experience. Perhaps Advanced Automation will indeed go its own sweet way, regardless of Cognitive Science; but if it does so, I fear that the resulting spin-off is more than likely to inflict multiple injuries on human society.

It is not, unfortunately, clear what injuries Longuet-Higgins had in mind.

Donald Michie's response to the Lighthill report was the most critical, substantial and wide-ranging. Before examining it, we need to know a little more about his background. Michie had worked with Alan Turing at Bletchley Park, and shared the fascination in the question whether machines could think. With 'no opportunities for him to follow up these interests after the war', except on the hobbyist level, Michie 'took medical sciences at Oxford and subsequently specialised in genetics'.³⁸ (I don't know whether this lack of opportunities related to the communism he shared with his wife,

³⁸ Donald Michie, *On Machine Intelligence*, Edinburgh: Edinburgh University Press, 1974. James Fleck, 'Development and establishment in artificial intelligence', in Norbert Elias, Herminio Martins and Richard Whitley (eds.), *Scientific Establishments and Hierarchies, Sociology of the Sciences* (1982) 6, pp. 169-217, p. 183.

the accomplished geneticist Anne McLaren.³⁹) He was appointed as a Reader in the Department of Surgical Science in Edinburgh in 1958. A visit to the United States in 1962 opened his eyes to AI developments there. Returning to the UK he began lobbying for better computer facilities. In 1963 he – remarkably - moved out of the Department of Surgical Sciences and set up his own ‘unofficial unit’, the Experimental Programming Unit. His lobbying of the government, specifically the DSIR, paid off, and he began to attract large grants, as well as organising a series of Machine Intelligence Workshops. With entrepreneurial skill, Michie by 1966 not only had secured official recognition from the university, in the form of the Department of Machine Intelligence and Perception, but he also pulled in the talents of Richard Gregory (from Cambridge, psychology of perception) and Longuet-Higgins (from Oxford, a very bright theoretical chemist, looking for a new challenge), with the intent of collaborating on building an intelligent robot. Michie also built up an international ‘Firbush’ network of supporters and collaborators in machine intelligence. Taking its name from Firbush Point on Loch Tay, where Michie hosted a three day meeting of the network, the Firbush newsletters convey a strong sense of excitement over the ambitious research underway by the early 1970s.⁴⁰

Michie had seen a leaked March 1972 version of the Lighthill report (minus two pages, warned his informant, that were ‘not being circulated to anyone, since they contain detailed comments on individuals’ – almost certainly about Michie himself).⁴¹ His public response, however, came as part of the published paper symposium in 1973. First, unlike the other respondents, he attacked Lighthill’s

³⁹ Alex May, ‘Dane Anne Laura Dorinthea McLaren (1927-2007)’, *Oxford Dictionary of National Biography* <https://doi.org/10.1093/ref:odnb/98949>. Anne McLaren and Donald Michie married in 1952, divorced in 1959 but remained close friends. Donald Michie resigned from the Communist Party of Great Britain in 1956, while McLaren remained a member.

⁴⁰ The meeting at Firbush Point, Loch Tay, involved two British (both Edinburgh, Michie’s Department and Meltzer’s Department of Computational Logic), one Russian (Mathematical Institute of the USSR Academy of Sciences, Novosibirsk), and four American bodies (SRI, MIT, Syracuse, and ARPA). *Firbush News* (June 1972) 2, p. 4, reproduced a letter from Lighthill.

⁴¹ British Library Add MS 88958/1/384. Copy of first 35 pages of Flowers’ submission to the SRC, sent to Michie by ‘J’, who had been given a copy by Freddy Lines (presumably Alfred Lines who had worked, like Lighthill, at RAE).

categorisation directly, which he characterised as ‘remote’, ‘misleading’ and especially unfair on category B’s work’s supposed foundations:

Most people in AI who have read the report have had the feeling that the above classification is misleading. Sir James has arrived at his position by interpreting AI as consisting merely of outgrowths from a number of established areas, viz.:

A as an outgrowth from control theory,

B as an outgrowth from science fiction,

C as an outgrowth from neurobiology,

These interpretations are remote from those current in the field itself.⁴²

Second, he argued that Lighthill had not consulted widely enough, in particular he had not invited opinions of the leading American AI workers: John McCarthy, Marvin Minsky, Nils Nilsson, Bertram Raphael and Alan Robinson, for example. Third, he attacked category B directly. His initial reading, written out in the first heat of reading the report, had been to suggest B be best understood as a ‘channel of communication between workers occupying the two poles [of A and C]’, a visualisation he had previously used as early as 1964.⁴³ Now, however, Michie chose to double down on the importance of category B work. Whereas Lighthill had labelled B merely Building Robots, Michie

⁴² Donald Michie, ‘Comments on the Lighthill Report and the Sutherland Reply’, in SRC, *Artificial Intelligence: a Paper Symposium*, op. cit. Michie would continue to complain that Lighthill had been unqualified, too hasty, distracted, and unwilling to consult AI experts widely enough (especially outside the UK), well into the 1980s. See: ‘Artificial Intelligence research’, prepared for submission to the House of Lords Select Committee on Science and Technology and drafted as a letter to Nature, 12 January 1980. British Library Add MS 88958/1/382.

⁴³ British Library Add MS 88958/1/385. Michie, ‘On first looking into Lighthill’s “Artificial Intelligence”’, August 1972. Michie own report on the field, delivered to the SRC, *Computing Science in 1964: a Pilot Study of the State of University-based Research in the United Kingdom*. London: Science Research Council, 1965, had represented work as a “dumbbell”, with the bells A (‘Advanced Computer Science and Automation’) and C (Computer-based CNS Research) linked by a connecting piece – ‘or bridge’ – marked with a question mark and seen as a ‘channel of communication’. Michie placed ‘Machine Intelligence’ in bell A, near the bridge, and ‘Computational Psychology’ in bell B, also near the bridge. The two pictures of the field – Michie’s in 1964 and Lighthill’s in 1972 are so similar that Lighthill must have, directly or indirectly, drawn his map of artificial intelligence after Michie’s own visualisation – a remarkable irony.

argued it should have been called Intelligence Theory.⁴⁴ Reducing the ambition to articulate a full Intelligence Theory to Building Robots, was, he suggested, as ridiculous as reducing the investigations into flight to merely 'Building Wind-tunnels', when the 'true bridge' the equivalent of basic intelligence theory would be the prestigious and successful science of aerodynamics. (No coincidence, I think, that aerodynamics was one of Lighthill's areas of expertise.) 'The equivalent science in the case of AI is at a primitive stage', said Michie, but 'It is the hope of every AI professional to contribute in some way to bringing the required theory into being'. Fourth, despite this aspiration towards theory, Michie, claimed category B work did indeed have 'practical benefits', both short and long term:

The subject, in so far as it comes within the Computing Science Committee's realm of interest, is concerned with machines, and in particular computers, displaying characteristics which would be identified in a human being as intelligent behaviour. Perhaps the characteristics which are most important are those of learning and problem solving. The applied benefits which may be gained from work in this field could bring considerable economic benefit to the country. They are two-fold:

- a. To relieve the burden at present on the systems analyst and programmer in implementing applications;
- b. To enable new and more complex applications to be undertaken in this country in competition with work elsewhere.⁴⁵

Finally, he asked for further support from the SRC, notably the supply of small PDP 10 machines, on which many American AI programs ran. With these arguments faltering, Michie tried other tactics, such as arguing that machine intelligence research was relatively inexpensive and potentially

⁴⁴ Michie, 'Comments on the Lighthill Report and the Sutherland Reply', *Artificial Intelligence: a Paper Symposium*.

⁴⁵ Some of the problems to be solved were discussed in Donald Michie, 'Machines and the theory of intelligence', *Nature* (23 February 1973) 241, pp. 507-512.

economically profitable (compared to nuclear physics or Concorde, say), attributing the mistake to a British preoccupation for counting pennies but wasting pounds.⁴⁶ Later, as the controversy opened up, John McCarthy and Richard Gregory joined Michie as respondents to Lighthill in a televised version of the paper symposium, filmed at the Royal Institution and broadcast as part of the *Controversy* series.⁴⁷

After the Lighthill Report the SRC reorganised the funding of affected fields, funding A and C separately leaving B in the lurch. This change in funding had dramatic effects, triggering a first AI winter and perhaps most directly impacting on Michie in particular. It came at a bad time for Edinburgh. Gregory had never settled and left for Bristol in 1970, while Longuet-Higgins and Michie had fallen out. Other departments were jealous of Michie's unit's freedom to research.⁴⁸ Even a writing campaign by Michie's American allies in artificial intelligence – Michie's Fimbush network – could not convince the SRC to change its mind.⁴⁹ With the robotics work attacked, Michie was essentially frozen out, with a new Department of AI being set up in 1974 without Michie, who was left with his own, small independent Machine Intelligence Research Unit.

Interpreting the Lighthill Report and Responses

I will offer a few specific comments on the Lighthill review process, before turning to my main argument. First, the extraordinary tone and focus of the report is surely facilitated by the fact of it

⁴⁶ Donald Michie, 'Machine intelligence in the bicycle shed', *New Scientist* (22 February 1973), pp. 422-423.

⁴⁷ See YouTube and also via Edinburgh's current Artificial Intelligence Applications Institute website <http://www.aiai.ed.ac.uk/events/lighthill1973/>. McCarthy's response can be found in the review of the paper symposium, published in *Artificial Intelligence* (1974) 5, pp. 317-322. 'If we take the [ABC categorisation] seriously', wrote McCarthy, 'then most AI researchers lose intellectual contact with Lighthill immediately'.

⁴⁸ 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'. Fleck, op. cit., p. 190.

⁴⁹ British Library Add MS 88958/1/383. Charles A. Rosen (SRI) to Flowers, 1 June 1973. British Library Add MS 88958/1/384. Nils J. Nilsson and ten others (from SRI, Stanford, MIT and University of Southern California) to Flowers, 16 May 1973. Rutovitz to Edwards, December 1973, proposes the example of a planned Jet Propulsion Laboratory Mars "crawler" as a practical case of category B work.

being a single-authored report. Lighthill did not have to moderate his arguments or make concessions in negotiating a text. Second, the incisiveness might have been anticipated and even expected by the SRC. Its Engineering Board's Computing Science Committee's long-range panel, completed earlier but published in June 1972, had been considerably more positive and yet had not resolved doubts that were clearly held about some work in the UK classed as AI.⁵⁰ Recall that Flowers, in his first letter to Lighthill, had given something of a steer to the applied mathematician when he talked of 'plausible charlatans'. Third, the format of the paper symposium, while allowing the targets of Lighthill (as well as allies) to air their views, was also invidious. Sutherland at Sussex for example felt encouraged to articulate criticisms of the work at Edinburgh.⁵¹ Given the platform and surely sensing danger, Longuet-Higgins had fashioned an alternative 'cognitive science' frame that both defended a subset of research while distancing it from the targeted Category B work.

Fourth, in the best sociological account of the history of AI in the UK, James Fleck accounts for the Lighthill report in three ways. He argues that Lighthill was acting conservatively, an 'affirmation of the status quo'.⁵² Specifically, Lighthill, says Fleck, 'affirmed the value of currently existing areas' with strong links to existing disciplines (A with control engineering, C with neurophysiology). Fleck also finds a 'prestige hierarchy' in operation, with 'those closest to the physical sciences accorded most prestige'. That analysis sort of makes sense of work in category A, but for category C Fleck has to rather vaguely suggest Lighthill must have been influenced by 'the dominance in [continental] Europe of the neurosciences ... over the cognitive sciences'.⁵³ Finally, Fleck says that Lighthill's response was merely typical of humanistic responses to AI, 'one among a multitude of attacks on the field'. The response here was therefore one of those identified by Turing when he had discussed one

⁵⁰ 'Report of Long-Range Computing Research Panel', Appendix 1 of *Policy and Programme Review of Research and Training in Computing Science*, London: Science Research Council, 1972.

⁵¹ 'The one postgraduate course (at Edinburgh) places too much emphasis on conventional mathematics and logic and does not capture the essence of the subject as understood in the US.'

⁵² Fleck, *op. cit.*, p. 209, p. 205.

⁵³ Fleck, *op. cit.*, p. 205.

of the main reasons we wrongly reject the idea of machines that think, that it challenges our supposedly special humanity.⁵⁴

Lighthill's report was taken at the time by many commentators, and has certainly passed into disciplinary lore in the way the story has been told, as an attack on AI. It has been credited with not only damaging AI work in the UK for a decade (not least in Edinburgh) but with such international repercussions that it was subsequently classed as the cause of the first AI winter. But, my argument is that if we look closer, and look in particular at Lighthill's wider vision of the proper relationship of research to practical needs, another interpretation emerges.

Lighthill's Cause

Recall that, to summarise Lighthill, category A was praised because it linked to industry and solving industrial problems, category C was praised for its engagement with scientific problems of the nervous system (which also linked to the practices of medicine and psychiatry), while the failing bridge, category B, did neither. Category B would continue to fail because it had no 'field of application'.⁵⁵ Further support for this interpretation of an underlying position on the application of science behind the Lighthill Report can be found. Indeed, I want to suggest that a commitment – even a principle – that science flourishes when engaged with practical problem-solving, was a thread that goes through much of Lighthill's work. It is therefore no surprise that he should view AI through this lens.

⁵⁴ Alan Turing, 'Computing machinery and intelligence', *Mind* (1950) 49, pp. 433-460.

⁵⁵ Note however that Fleck refers in passing to Michie's projects as actually being entrepreneurially commercial: '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'. Did Lighthill know or ignore or set aside this applicability?

For example, in 1962 Lighthill gave a speech to the members of the Chemistry Department and the Metallurgy and Physics Department at the Royal Aircraft Establishment. He spoke of his 'deep conviction of the vital importance of what [he called] ... "applied high-grade science"', by which he meant applied research with a long-term view. His exemplar was Bell Labs. He stressed the economic importance of this work solving the problems leading to advanced technologies (citing the Prime Minister, Harold Macmillan, in support). He then addressed the relationship between research judged on scientific merits and research that addressed such problem-solving, using a metaphor that itself came straight from applied mathematics:

I think it is unfortunate that academic research has ... failed to take this [focus] into account. Sometimes it has taken the view that nothing matters except the production of high-grade scientific work. However, if you apply this extra boundary condition, that one has selected the work to be done because of its potential for advanced technological application, or allowed the progress of the research to be guided by the possibilities that one can see, and if under this extra boundary condition you still manage to do high-grade scientific work, you have evidently acquired even greater merit.⁵⁶

In other words, scientific work of the highest merit was that framed by practical problem-solving for industrial benefit within which scientific excellence was pursued and achieved. We saw the same thought behind Lighthill's suggestion of redirecting UK fusion research to industrial problem-solving. We see it again in Lighthill's promotion of applied mathematics, both as subject and in the form of new institutions he thought necessary.

⁵⁶ MJL D8. Lighthill, 'Transcript of the speech made by the Director to members of Chemistry Department and Metallurgy & Physics Department at the Assembly Hall on the 1st November 1962', 1962. He uses the boundary condition metaphor again in MJL K14. Lighthill, 'Synopsis of speech given at the Fourth British Theoretical Mechanics Colloquium, Bristol, 1962', 14 March 1963.

In 1962 Lighthill gave a speech at the Fourth British Theoretical Mechanics Colloquium, held in Bristol. He surveyed the subject matter of applied mathematics and its relationships to neighbouring sciences:

The good applied mathematician must understand a physical situation ... as well as the mathematical constructions that he uses to illuminate it. From his knowledge of the neighbouring sciences ... he creates a whole new world of new relationships between them. Without physics, he lacks depth; without [pure] mathematics, he lacks power; without engineering, his work lacks practical value.⁵⁷

The main subject matter of applied mathematicians, said Lighthill, was 'above all, the art of *bridging* the gulf between the mathematical and the physical; between, that is, the world of numerico-logical constructions and the world of experimental observations'.⁵⁸ And for this bridging work to be successful it must, said Lighthill, be responding to practical problems.⁵⁹ He ended the speech by calling for a new institution: we must 'consider whether the time has come for the needs of applied mathematicians to be served by ... a new professional body'.

Within a couple of years a new Institute of Mathematics and its Applications had been set up, with members of different grades, a lecture programme, two new journals and larger meetings planned.⁶⁰ The Leverhulme Foundation provided the funds.⁶¹ By 1965 the IMA had 392 fellows, 174 associate fellows, 40 companion members, 70 graduate members and 6 student members.⁶² Lighthill was chair

⁵⁷ MJL K14. Lighthill, 'Synopsis of speech given at the Fourth British Theoretical Mechanics Colloquium, Bristol, 1962', 14 March 1963.

⁵⁸ My emphasis.

⁵⁹ 'One essential thing to a successful applied mathematics school is closeness of contact of at least some of its members with areas where the work is successfully applied'. 'Summing up, the applied mathematician to-day has to span the whole quadrilateral [rigorous mathematical reasoning – full numerical computation – equivalent physical argument – application to practical problem] and to forge clearly, also, all the links I have shown here' (see figure).

⁶⁰ The articles of association are in MJL K14. The first registered address was 29 Gordon Square (Gordon House). Its working office was in Maitland House, Warrior Square, Southend-on-Sea, Essex.

⁶¹ MJL K28. Sir William Hodge (Joint Mathematical Council) and Lighthill, Draft funding application to Leverhulme Foundation, undated (1964). MJL K30. Lighthill to Pack, 11 February 1965

⁶² MJL K29. J.T. Cambridge to Lighthill, 1 January 1965.

of the Provisional Council and its first President. A description of the ideal applied mathematician can be found in Lighthill's speech on the occasion of Sir Geoffrey Taylor being made the first honorary fellow of the IMA. He drew attention to how Taylor's ground-breaking theoretical work was 'most intimately linked with, and usually suggested by observation, and subjected to the rigorous test of experiment'; furthermore, he 'sought direct personal experience of every subject studied', whether it was gathering the first statistical knowledge of turbulence from observations of flying kites off the Grand Banks of Newfoundland, or designing ways to secure the Mulberry Harbour from experience with sailing.⁶³ Again we find the combination of practical problem solving and high-grade mathematics being praised.

In 1977, Lighthill was lecturing the IMA on 'Bridging the chasm separating examination questions (even "difficult" ones) from real-world problems (even "easy" ones)', showing that he wanted to apply his guiding principles right into the pedagogical techniques and measures of his subject.⁶⁴ The 'real world' here encompassed 'industry, commerce and government'. 'I recommend for mathematics undergraduates', Lighthill told his audience, 'the availability of PROJECT WORK IN DEPTH ON REAL-WORLD PRACTICAL PROBLEMS'.⁶⁵

In 1982, Lighthill was contacted again by the research council (now the Science and Engineering Research Council, SERC). John Kingman – another high-powered applied mathematician, this time of statistics rather than of fluid mechanics – informed him that a new report on AI had been commissioned, with Roger Needham and Peter Swinnerton-Dyer as the authors, whose central conclusion was that 'research in this area has become less "millennial" and philosophical, and this of course is a tendency which SERC would wish to encourage'.⁶⁶ They supported the proposal

⁶³ MJL K36. 'Remarks by the President on presenting a certificate of election to the Inst. First Honorary Fellow. Sir Geoffrey Taylor', undated.

⁶⁴ MJL J35. Lighthill, 'Bridging the chasm separating examination questions (even "difficult" ones) from real-world problems (even "easy" ones)', lecture delivered 4 May 1977.

⁶⁵ His capitalisation. Lighthill and others published a book with Penguin *Newer Uses of Mathematics* (1977) that gave examples.

⁶⁶ K685. Kingman to Lighthill, 21 October 1982.

recommended by the Alvey committee for a 10 year industrial strategy of extensive and generous funding of intelligent knowledge based systems (IKBS).⁶⁷ They also commented on the Lighthill report:

It must be recognised that the Lighthill Report was right for its day, but that its day is past. AI is a different subject from what it was ten years ago, and it is set in a different context. Then it was the wild blue yonder. Now it is the open frontier of IT...⁶⁸

What, asked Kingman, did Lighthill make of it? Lighthill certainly didn't object to being told the days of his report were past. Instead he was full of praise for the Alvey programme (and note why):

...the Alvey recommendations for major R&D and educational initiatives in the field of Intelligent Knowledge Based Systems seem to me to be exactly right. They are combined, as they should be, with major proposals in three other, equally vital, technological areas. They are *directly linked to industry and its needs*.⁶⁹

Furthermore, he revealingly reflected on the Lighthill Report, and his words need to be quoted at length:

... the Alvey recommendations steer clear of a certain dangerous philosophical misconception which was seriously impeding progress in 1972. This misconception may have had what you call a "millennial" element in it, but primarily it was just wrong-headed. I am speaking, of course, of the nineteen-sixties view that progress in AI would come out of viewing it as a seamless robe stretching all the way from advanced automation to neurobiology. Progress was greatly facilitated when that view had largely been abandoned.

⁶⁷ K685. Needham and Swinnerton-Dyer, 'Artificial Intelligence research in the UK', October 1982.

⁶⁸ Needham and Swinnerton-Dyer made a number of recommendations, one of which was for the strengthening of Edinburgh, a 'significant upgrading of its staff and facilities'.

⁶⁹ K685. Lighthill to Kingman, 17 November 1982. My emphasis.

Neurobiology flourished amazingly, and so did Computer Science; but in completely different ways.

Such a change in approach ... had been necessary in Britain's most important AI centre, and I have marked this letter PERSONAL & CONFIDENTIAL because I need to mention individuals at this stage. It was, above all, the work at Edinburgh which ... was being seriously held back by the approach seeking to make links towards neurobiology *instead of towards practical needs*. The direction of the work was distorted through the strongly held view of Donald Michie (who had nailed his colours to this particular philosophical mast); simultaneously, his personal characteristics tended to foment additional, very serious difficulties...⁷⁰

The Alvey programme, which linked academic research to industrial (and military) needs, and was well funded (£350 million), was precisely what Lighthill took to be the appropriate relationship of research to the wider world. In contrast, and to further support this interpretation, he criticised Needham and Swinnerton-Dyer for backsliding on this point: they were 'trying to pull the subject back from a *practical, problem-oriented* but intellectually demanding outlook into a philosophical stance...'.⁷¹

Pro- and Anti-Working Worlds

⁷⁰ My emphasis. The difficulties were the reason for the departure of Richard Gregory and Christopher Longuet-Higgins, Lighthill asserted. He went on to say that 'the change of direction and of organisational structure that followed my Report greatly helped AI work at Edinburgh to flourish. You may care to ask John Burnett, the Principal, for his view on this (the view of the former Principal, Michael Swann, was identical). He will tell you how the atmosphere engendered by my Report allowed the University to adopt arrangements with the main work in AI proceeding very satisfactorily while Michie moved to the sidelines (and, this year, he accepted the University's offer of early retirement)'. For the creative destruction of interdisciplinarity, specifically, the destruction of the bridge between biological and computer scientists, see also: Lighthill, "The development of AI has been very much as I forecast". Interview with Sir James Lighthill, *Machine Intelligence News* (1986) 2(5), p. 3-4.

⁷¹ My italics but I have no doubt that this is Lighthill's emphasis. He continued '...oriented towards neurobiology but failing to recognise its biological and biochemical complexity. They want AI to dabble in the neurosciences, which are prospering, at the expense of British industry, which is not!'

Running through Lighthill's career was therefore a principle: that the very best scientific work consisted of a tight relationship between research and practical problem-solving. In addition, such work was of the greatest national importance. Furthermore, again this relationship was conceived as a bridge. The Lighthill Report on Artificial Intelligence was not a specific attack on AI but a review of AI guided by this principle. Seen in this light, we can see why Lighthill praised the parts of AI he did as well as, of course, the part – the mainly Edinburgh part – of AI that he heavily criticised. We can also see why category B – Building Robots but also the supposed Bridge between industrial automation of category A and the neuroscience and psychology of category C – was, for him, problematic. It was in two senses a failed bridge: it failed to link the two, and the reason it did was because it failed to aim to solve practical problems. Lighthill is therefore an example of a historical actor adopting what I call a 'pro-working worlds' stance: he considers the orientation of good science to be towards problem-solving and articulates what the responsibilities and responses of scientists should follow.

In contrast, Donald Michie, like most scientists, wrestled with the issue of presenting his programme as relevant to practical problems whilst at the same time preserving, through anti-working world argument, autonomy and space to pursue his research interests freely. Let's review two moments when he addressed these points. The first contains some "pro-working worlds" aspects, while the second is call for necessary independence from working world concerns. In his 'On first looking into Lighthill's "Artificial Intelligence"', Michie asked himself the question 'Is your robot really necessary?'.⁷² His answer illustrates his way of thinking about science's representatives:

If the "knowledge base" which a laboratory chooses to investigate is the one which underlies the ordinary physical transactions of everyday life ... then the knowledge must somehow, in

⁷² British Library Add MS 88958/1/385. Donald Michie, 'On first looking into Lighthill's "Artificial Intelligence"', August 1972.

boiled down form, be got into the machine. Such a boiled down form is referred to by the MIT workers as a “mini theory” ... Lighthill might refer to such a system as “nursery physics”.

So the “mini theory”, “nursery physics”, or what have been called “microworlds”, might be the abstract representative of the real world. But then ‘is it necessary actually to build a robot?’. Michie found an answer in a ‘crushing epigram’ he attributed to the cognitive psychologist R.L. Gregory: ‘the cheapest store of data about the real world is the real world’.⁷³ His point was that ‘the test of a given machine representation of knowledge about the real world is its ability to support effective interaction with it’, and to do so it must either form a ‘software simulation of a real world’ or be presented with a ‘fragment of the real world via a robotic peripheral’. Given ‘how much it would cost realistically to simulate even a fragment of the real world’, robots were the cheapest, best option. Furthermore, such representatives related to working worlds, since it was the ‘research worker’s wish that his chosen experimental tasks should foreshadow, even dimly, possible applications’. While this places Michie down the spectrum away from Lighthill’s extreme pro-working world position on applications, it clearly enabled him to justify his programme in terms of practical benefits. Indeed Edinburgh’s “hand-eye” robotics research, centred on ‘packing objects into containers’ and ‘assembling structures from descriptions’, opened the way, wrote Michie, to ‘computer-controlled packagers, assembly-line tools, dockside cranes and forklifts, building site bulldozers and mobile shovellers and grabs, not to mention wider software techniques of associated scheduling and control systems’.

Yet Michie was also able to articulate an anti-working worlds position. This view is clearly seen in a letter written to the Essex computer scientist Yorick Wilks in 1982. Wilks had asked for feedback on a paper outlining possible approaches to take during the Alvey programme on Intelligent Knowledge

⁷³ The original source is probably Edward Feigenbaum, ‘Artificial intelligence. Themes in the second decade’, Advanced Research Projects Agency (ARPA) Memo No. 67, 15 August 1968. Feigenbaum, quoting a view from within the SRI group, wrote that ‘the most economic and efficient store of information about the real world is the real world itself’.

Based Systems (IKBS), the first major revival of UK research council support of artificial intelligence, mentioned above. Specifically, Wilks asked whether there was a 'unified theory' of machine intelligence 'that was worth working on, independent of building some more working [intelligent systems] in the UK?'⁷⁴ Michie reached back into the history of science to formulate his reply:

Yes, in the same sense that a unified theory of mechanics was worth working on during the hundred years of pre-Newtonian development by the civil engineers (mainly Italian) and was worked on (Galileo's rolling cannon-ball experiments ... and Leonardo's earlier partly successful attempt to put together a system of "laws of mechanical forces").⁷⁵

But such a theory was not worth doing separate from building actual intelligent systems, 'any more that it would have paid off to work on the Chromosome Theory of Heredity independent of building the first genetic maps (of fruit fly chromosomes)'. Indeed Michie, returning to the Renaissance, cast himself as Galileo:

I cannot think otherwise, since the FREDDY series of [Edinburgh robot] projects were actually (little known fact) built for doing "cannon-ball experiments", a sort of epistemoscope you might say, and was only presented as manufacturing technology as a result of pressure from well-wishing SRC officers in a last ditch attempt to save it. ...

The situation really is a bit like renaissance mechanics and optics ... but one problem is that the [research council] looks over its shoulder continually at public justification in terms of industrial, military or medical pay-off. ...

In this version, expressed privately from one academic researcher to another, the orientation to practical problem solving was merely presentational. The actual motivation was the building of *epistemoscopes* – instruments for the discovery of new theory and knowledge:

⁷⁴ British Library Add MS 88958/1/385. Yorick Wilks and Bruce Anderson, untitled, 6 December 1982.

⁷⁵ British Library Add MS 88958/1/385. Michie to Wilks, 13 December 1982.

for what we are in a position to do (just as, to an extent, the atomic physicists of the thirties and forties, or the molecular biologists of the sixties and seventies) is to forge a new science – the most consequential there has ever been.

Therefore let [the research council's] public relations machine put out whatever emphases it judges to be politic within its world but let us build our “epistemoscopes” (however closely we may co-operate in this with industry etc) for the honest and honourable reason that motivates us as academic professionals.

Conclusion

This paper has had two aims, one narrow and one broad. The narrow aim was to inquire into the causes, motivations and content of the Lighthill report on artificial intelligence, which interrupted research in the field for a decade after its publication in 1973. I have argued that behind James Lighthill's criticisms of a central part of artificial intelligence was a principle he held throughout his career that the best research was tightly coupled to practical problem solving. We only see that the report was not a specific attack on AI when we place the report in the context of Lighthill's longer vision and career. It was also the case, however, that the Science Research Council did have narrow concerns about the variety and quality of artificial intelligence research, and, as the instruction from the SRC chair, Brian Flowers, reveals, certainly gave Lighthill a steer with respect to the target of the review. Neither of these aspects of the Lighthill report have been previously revealed.

The broader aim of the paper was to begin mapping how scientists (and in Lighthill's case, a mathematician) have articulated and justified relationships between research and problems. The 'working worlds' model offered an analyst's category of how such relationships operated based on a survey of twentieth century. But a deeper understanding will come through historicization, hence the project to uncover the range of articulations. Lighthill's principle represented a view of a strong,

tight relationship. In Donald Michie's response to the Lighthill report I have reported at least two other types, either as a (perhaps) cynical tying of research projects to practical aims in order to meet patron's expectations. But Michie also articulated a less cynical version, one in which the 'honest and honourable' motivation of the scientist was building of 'epistemoscopes', instruments whose purpose was not direct or immediate problem-solving but rather the revelation of unified theory. Lighthill's instruction to support "practical problem-oriented" mathematics and Michie's plea to be allowed to build pure, honest "epistemoscopes" in order to discover new theory illustrate two actor's category articulations of the relationships between science and working worlds.

So what? Why are studies such as mine above worth doing? What are the consequences of having such case studies? Perhaps it would help if I sketch an ultimate destination. I am very interested in finding and describing the conditions under which science can best engage with the world's problems. I want this picture to be empirically robust, built from the best of archival research, and justify recommendations for policy change. Even just in the UK we have moved from policies of active state intervention and support of emerging science-based technologies, through a period from the late 1980s when government funding for "near market" research was cut back, to a most recent phase when, again, having an active industrial strategy is championed. These shifts in policy, which often focussed on how the academic science base should relate to industry and wealth creation, paid little attention to how scientists actually understood the relationship of their work to problem solving. I have shown elsewhere that the ending of "near market" research was based on selective, anecdotal history of science, presented by a political adviser to prime minister. Here is where professional, historical, archival research is essential. Not only can the variety of actual stances towards problem-solving be uncovered, described and mapped, but also differences between, if you will, relatively on- and off-the-record stances, by comparing the arguments found in documents designed for wide circulation with more private correspondence. The scientist who might tell a public body such as a research council that their work aimed for industrial application might privately and honestly hold that their true purpose was theoretical revelation. A mathematician

might make a publicised attack on a single branch of science but it is better interpreted, after historical investigation, as stemming from a conviction that the best science comes from deep engagement with problem solving.