

A History of the Early Years of Artificial Intelligence at the University of Edinburgh

The European Journal on Artificial
Intelligence
2026, Vol. 39(3) 527–540
© The Author(s) 2026



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/30504554261417567
journals.sagepub.com/home/eai



Christopher KI Williams¹ , Vassilis Galanos²  and Xiao Yang¹ 

Abstract

This article describes the early years of artificial intelligence (AI) at the University of Edinburgh, roughly from the early 1960s to the mid-1980s. It covers the key founders and various administrative structures, research highlights and teaching developments from this period, notable women in early AI research, and discusses the national and international connections of the work carried out at Edinburgh. Motivated by a scarcity of historical documentation on this pivotal period, we employ a methodological blend of archival research and recent oral histories to draw a detailed picture. The study highlights the contributions of key figures such as Donald Michie and Christopher Longuet-Higgins, whose interdisciplinary work established Edinburgh as a European hub for AI research. Despite challenges like the “AI winter”, sparked by the Lighthill Report, Edinburgh fostered advances in many areas of AI. The analysis also explores early ethical considerations in AI development, reflecting on technological neutrality and the anticipation of AI’s societal impact. This historical account not only clarifies Edinburgh’s critical role in AI’s evolution but also offers enduring lessons on interdisciplinary collaboration and ethical responsibility, relevant for the continued advancement of AI technologies today.

Keywords

artificial intelligence, cognitive science, Edinburgh

Received: March 4, 2025; accepted: January 8, 2026

1 Introduction

This article was motivated by the celebrations of 60 years of artificial intelligence (AI) at the University of Edinburgh in 2023, and the relatively small quantity of available research into its history. We – an Emeritus Professor of Machine Learning, an early career Lecturer with expertise in the historical sociology of AI, and a doctoral candidate at the intersection of bioinformatics and the history of science – have joined forces to shed further light into the early foundations of AI at Edinburgh and extrapolate lessons that are relevant for the greater history of AI in Europe and beyond, as well as its future.

This article has the following structure. We first outline our methodology. We then discuss the key founders who shaped the early AI environment at Edinburgh, eventually leading to the establishment of specific institutional structures from informal groups to research units and University Departments or Schools. We offer a brief explication of challenges faced during the development of these structures, typically placed under the umbrella term “AI winter” – although we show that the reasons behind this period of AI disillusionment, within the Edinburgh context, go beyond the mere under-delivery of technical promises (as compared to other regions that faced “AI winters”), and involve, also, the social collision between ambitious agendas and scientific egoism. We then discuss technical research highlights and women in AI research. This is

¹School of Informatics, University of Edinburgh, Edinburgh, UK

²Lecturer in Digital Work, University of Stirling, UK

Corresponding Author:

Christopher KI Williams, School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, UK.

Email: c.k.i.williams@ed.ac.uk

followed by coverage of commercialization and teaching activities from the era. We also outline the national and international connections between Edinburgh and pivotal centres of research excellence. We finally report on available knowledge about AI's responsible and ethical development in this early era, opening up room for current scientific responsibility in the era of more advanced AI techniques and applications.

This article is a considerably expanded version of the workshop paper by Williams et al. (2024).

Methodology: Close study of existing historical accounts of AI in the UK, and Edinburgh in particular, gave an initial shape to the historical outline (Agar, 2020; Fleck, 1987; Howe, 2007) while works on specific or more general aspects of AI history assisted in situating Edinburgh's contributions in context (Crevier, 1993; Nilsson, 2010; Olazaran, 1996). This study has benefitted from previous engagement with the same corpus (Galanos, 2023). These resources were augmented by conducting new interviews in 2024, focusing on the oral history of AI at Edinburgh (as well as the neighbouring disciplines of cognitive science and computer science). These interviews are available at https://groups.inf.ed.ac.uk/aics_history/oral_history.html.

The more recent oral historical accounts, having the privilege to be sufficiently distant from the period of excessive change and challenge in the field of AI, allowed for a fresh and critical interrogation of the events. We have analysed the interview data employing iterative axial coding, initially based on our interview plan's thematic structure, however, allowing room for surprising new findings. In particular, some themes about the degree of intensity of controversy between key actors in Edinburgh's AI scene, as well early conversations about AI ethics, have emerged throughout the interview process, confirming our adoption of an abductive methodological framework (Timmermans & Tavory, 2012). To ensure rigorously and trustworthiness, we employed evaluation criteria for qualitative research as outlined by Lincoln and Guba (1986), including, but not limited to, external audits through a relevant conference presentation, peer debriefing with participants and colleagues sharing the same intellectual spaces with us, triangulation of historical resources and comparison with findings. Our research also benefits from prolonged engagement of the authors with the subject area.

2 Founders, Structures and Challenges

2.1 Beginnings

We first provide an overview of the various structures in the University of Edinburgh that were developed to accommodate AI research in Edinburgh from the early 1960s to mid-1980s, and the people involved in their foundation. This narrative involves the development of Schools, Departments and research units, often echoing the research directions envisaged by influential individual researchers. Our reconstruction of the timeline of the various structures created is shown in Figure 1.

The start of AI research at Edinburgh can be traced back to a small research group established at 4 Hope Park Square in 1963 under the leadership of Donald Michie. In 1965 this became recognised by the University as the Experimental Programming Unit, with Michie as Director. Michie had joined the University in 1958 as a Senior Lecturer in Surgical Science, but he had also worked from 1942 to 1945 at Bletchley Park as a codebreaker. Copeland (2023) provides a short transcript of Michie describing conversations with Alan Turing about AI topics during wartime at Bletchley Park. After World War II Michie was keen to pursue "the possibility of constructing machines that could think" (Fleck, 1987, p. 121), but at that time there were no opportunities to further this goal. In 1953 he completed a DPhil in mammalian genetics (Lee, 2012). He was a gifted communicator, not just in the scientific literature but also in the popular press; for example he was for many years in the 1950s the Daily Worker's science correspondent.

By 1961 Michie had already published some research on reinforcement learning (Michie, 1961). In 1962 he visited the United States and became aware of the developments in AI and the computer facilities available there. On his return he set up his research group, and lobbied energetically for improved computer facilities in the UK.

In 1966 the Department of Machine Intelligence and Perception (DMIP) was formed, funded by a large Science Research Council (SRC) grant held by Donald Michie, Christopher Longuet-Higgins and Richard Gregory. Longuet-Higgins had been an eminent theoretical chemist before switching to the study of the mind, having been appointed John Humphrey Plummer Professor of Theoretical Chemistry at Cambridge in 1954 at the age of thirty (Gregory & Murrell, 2006). Gregory had been a Lecturer in the Department of Experimental Psychology at Cambridge before he moved to Edinburgh.¹ Longuet-Higgins and Gregory had been looking to set up a brain research institute, and considered Cambridge, Sussex and Edinburgh as locations; Edinburgh came up with the most attractive offer, so Longuet-Higgins and Gregory "packed [their] bags for Edinburgh [...] in the shared conviction that the workings of the mind could not possibly be as tedious as the psychologists made them out to be, or as peripheral as the physical scientists tended to assume" (Longuet-Higgins, 1987, pp. xi–xii), and were welcomed by Michie, among others, to Edinburgh. Michie also introduced Longuet-Higgins to computing (Longuet-Higgins, 1987, pp. xii). The establishment of DMIP was supported by Michael

1963	Experimental Programming Unit (EPU) founded (Michie)
1966 - 1971	Metamathematics Unit (Meltzer)
1966 - 1971	Department of Machine Intelligence and Perception, with groups EPU (Michie), Bionics Research Laboratory (BRL) (Gregory), Theoretical Section (Longuet-Higgins)
1969 - 1985	School of Epistemics (founder: Longuet-Higgins)
1971 - 1974	School of AI, with sections Department of Machine Intelligence (Michie), BRL (Howe), Theoretical Psychology Unit (Longuet-Higgins), Department of Computational Logic (Meltzer).
1974 - 1998	Department of Artificial Intelligence (first head: Meltzer)
1974 - 1984	Machine Intelligence Research Unit (Michie)
1983 - 2001	Artificial Intelligence Applications Institute (AIAI)
1985 - 1998	Centre for Cognitive Science (CCS)
1998 - present	Division of Informatics (later School of Informatics) formed, including DAI and CCS

Figure 1. Approximate time line of units, departments and schools related to AI.

(later Lord) Swann, who became Principal of the University in 1966, and provided chairs for the new senior people (Fleck, 1987, p. 124). Michie was appointed to a Personal Chair of Machine Intelligence in 1967.

Michie's main interests were in design principles for the construction of intelligent robots, whereas Gregory and Longuet-Higgins were interested in using computational modelling of human cognitive processes to provide insights into their nature. Longuet-Higgins' research group was called the Theoretical Section, based at 9 Hope Park Square and later 2 Buccleuch Place. Gregory's group was called the Bionics Research Laboratory, and was based at 5 Forrest Hill. When Gregory moved to Bristol in 1970, Jim Howe took over leadership of the Bionics Research Laboratory.

Bernard Meltzer was another AI pioneer at the University of Edinburgh. He joined the university in 1955 as a Lecturer in the Department of Electrical Engineering. Following an interest in mathematical logic, he carried out exploratory work on programming computers to do mathematical reasoning (as well as research on solid-state and high-vacuum electronics). In 1964–65 he further explored this topic with an appointment at the SRC's Atlas Computer Laboratory, and he subsequently founded the Metamathematics Unit in the Faculty of Science in 1966.² In 1970 he was the founding editor of the *Artificial Intelligence Journal*, which became a leading journal in the field. In 1972 he was appointed to a Personal Chair in Computational Logic. In this era Edinburgh was one of the few centres in the world working on AI, along with Stanford, MIT and CMU. Michie, Longuet-Higgins, Gregory and Meltzer may be regarded as some of the "founding fathers" of AI.

In 1969 Longuet-Higgins founded the School of Epistemics, an interdisciplinary group which brought together people with an interest in the mind. Longuet-Higgins defined epistemics as "the construction of formal models of the processes (perceptual, intellectual, and linguistic) by which knowledge and understanding are achieved and communicated" (Longuet-Higgins, 1977). In 1974 when Longuet-Higgins moved to Sussex, Barry Richards of the Department of Philosophy became the Director of the School.

The value of the School of Epistemics was found in its facilitation of debate across disciplines. Minutes of the steering committee meetings of the School from the early 1970s indicate that there was representation from across a wide spectrum of disciplines, including many senior Professors: Educational Sciences (Prof L Hudson); English Language (Mr J P Thorne, who was also Academic Secretary to the School); Genetics (Prof C H Waddington); Linguistics (Prof J Lyons); Molecular Biology (Prof M R Pollock); Philosophy (Prof W H Walsh); Psychiatry (Dr N Kreitman); Psychology (Prof D M Vowles); Science Studies Unit (Dr D O Edge); MRC Speech and Communication Research Unit (Prof R C Oldfield); Zoology (Dr A W G Manning); as well as Drs Burstall and Meltzer and Prof Longuet-Higgins from the School of AI.

According to interviews, three important factors ensured that Edinburgh became a centre of excellence for AI in Europe. Firstly, the fact that the DMIP was unique within the UK in receiving substantial financial support in the late 1960s and early 1970s:

Well, I think it was pivotal for AI because it was the only well-funded AI centre in Britain. The Science Research Council, or whatever it was called then, decided to fund one big centre of AI. They wanted to try and keep up with the Americans, I guess. And they decided to do it in Edinburgh. (Hinton, 2024).

Secondly, the physical proximity between the different research hubs – the Experimental Programming Unit, the Theoretical Section, and the Bionics Research Laboratory – enabled the facilitation of important conversations and collaborations:

Hope Park Square physically is about 100 yards from Buccleuch Place where the Epistemics people were. So people who don't know the geography in Edinburgh, we're all still pretty much close together. And even Forrest Hill, which is only a quarter a mile from that, that's where the robotics originally was. (Tate, 2024).

Thirdly, the immense amount of variety in AI-related research concentrated within this very small area:

And at that time [the early 1970s], there were lots of other people working in AI, but they weren't really the world leading centres at the time, I believe that's true. So you've got Edinburgh as really the European Centre for AI, that's definitely the case anyway. So, it was an attraction, from that point of view. Donald Michie had a lot of interesting things going on. With robotics work on the Freddy robots as well at the time. But there's natural language work, there's work in search, there's work in a range of different topics. So it means there was a mix of things going on, and quite a few people. It was a centre of excellence, really. So that's what attracted me to come up. (Tate, 2024).

2.2 The Lighthill Report and Beyond

Growing concerns about the feasibility of promised outcomes, paired with rumours circulating about the integrity of the AI research culture at Edinburgh (as discussed in Agar (2020)) led the SRC chaired by Sir Brian Flowers to commission the eminent applied mathematician Sir James Lighthill to conduct research on the state of AI research in the UK, with particular focus at Edinburgh (Agar, 2020; Howe, 2007; Nilsson, 2010, p. 263). The result was the report "Artificial Intelligence: A General Survey" (Lighthill, 1973) in which Lighthill offered his outline of the AI field under the mnemonic "ABC", being the combination of (A) advanced automation, (C) computer-based central nervous system (CNS) research, and (B) building robots. "B" also stood for bridge activity, including hand-eye coordination, visual scene analysis, the use of natural language, and common-sense problem solving. Lighthill heavily criticised the future of these fields as a potentially unified whole due to perceived failures, especially in category B, stating: "By contrast, the sense of discouragement about the intended Bridge Activity of category B, centred upon Building Robots, seems altogether more widespread and profound, and this raises doubts about whether the whole concept of AI as an integrated field of research is a valid one." He also suggested that AI was a long way from being an applied domain suitable for industrial adoption, especially in category B. The effects of the report were widespread and damaging, and led to a loss of funding, in Edinburgh, for robotics work (Birse, 1994, p. 177).

The Lighthill report, and the year 1973, also marked the birth of (or at least the coinage of the term) cognitive science at Edinburgh, through a response to Lighthill by Longuet-Higgins, in defence of AI's contributions to science, in particular cognitive science (Longuet-Higgins, 1973) (therefore, not only defending his own version of AI, but also his own field, as per Agar's remark; Agar, 2020, p. 298).

Longuet-Higgins and Lighthill had been contemporaries at Winchester College, a leading English public school (Gregory & Murrell, 2006, p. 151). Hinton suggests that Lighthill's support for category C could have been boosted by their long-standing friendship since their schooldays at Winchester and their mutual membership of the Royal Society:

[...] that meant that Lighthill and Christopher, they're both in the same club – the Royal Society. Had known each other since high school days. And I think they respected each other. And so one interesting thing about the Lighthill report is there's a little out in it, which basically says AI is rubbish, unless you study how the brain works, in which case it's okay. And that sort of let Christopher out. (Hinton, 2024).

Lesser known aspects of the Lighthill debate include Agar's archival documentation of SRC chair Brian Flowers' directed steer towards Lighthill's assessment. In correspondence with Lighthill when initially contacting him to conduct the assessment, Flowers writes: "These subjects, and artificial intelligence itself, are highly complex, very much the preserve of experts and perhaps sometimes of plausible charlatans", eventually asking Lighthill "to make a considered appreciation of the subject of artificial intelligence, its achievements, its practitioners, its promise and its needs" (Flowers, cited in Agar (2020, p. 294)).

Other interpretations of the motives behind the Lighthill report place Michie's ambition towards advancing expensive robotics research and his rivalry with Longuet-Higgins at the epicentre of the investigation, or, as Fleck puts it, "irreconcilable differences between him [Longuet-Higgins] and Michie over the installation of the new computer and the robot project, as well as their approaches to work in the area" (Fleck, 1987, p. 126).

Michie was a very idiosyncratic character. He was not afraid to go right to the top for funding. As Gordon Plotkin recalls (unpublished interview, 2024) "when Donald Michie wanted something, he just went straight to the Principal and asked for it"

David Willshaw's insight into the Lighthill report highlights a critical moment where Michie's ambitious funding requests may have backfired: "if he had been more modest, he would have probably got the grant" (Willshaw, 2024). Willshaw also comments on tensions between Michie and Longuet-Higgins in the late 1960s/early 1970s:

[...] there were disagreements probably about where the Department was going, because Donald had his big robot project. And Christopher was really happy for him to be doing that. But he didn't particularly want to be involved with that himself, because he had his own agenda. (Willshaw, 2024).

Nevertheless, interpretations of the so-called "AI winter" differ across researchers from that time. For some, AI research never halted, it just continued under different branding based on the different specialisms within AI. AI researchers would continue to attend AI conferences (that had dramatically grown in size during the "AI winter") but would simply avoid the term "AI" in their applications for funding (Haigh, 2023). For others, it was indeed a time of research grief, and led "to the emigration of eminent British AI workers to the United States" (Crevier, 1993, p. 117). A recollection by J. Strother Moore confirms this: "We left Edinburgh in December, 1973, in part because of the Lighthill Report and the coming 'AI winter' in the UK. But the ideas we developed in the Metamathematics Unit lived on and had real impact" (Moore, 2021).

In 1971, in response to the disharmony in the department, there was a structural reorganization and DMIP was transformed into the School of Artificial Intelligence, run by a steering committee. Three of its research groups were the Bionics Research Laboratory (led by Jim Howe), the Department of Machine Intelligence (led by Donald Michie) and the Theoretical Psychology Unit (led by Christopher Longuet-Higgins). At the same time Meltzer's Metamathematics Unit joined the School as the Department of Computational Logic (Howe, 2007). Following the Lighthill report there was an internal review of AI by a committee appointed by the University Court, leading to the formation of the Department of Artificial Intelligence (DAI) in 1974. Its first head was Meltzer, who stepped down in 1977 and was replaced by Jim Howe, who led it until 1996. A separate unit, the Machine Intelligence Research Unit, was set up in 1974 to accommodate Michie's work. He remained there until 1984, when he left Edinburgh to co-found the Turing Institute in Glasgow. Despite the Lighthill report's influence, Michie continued to attract stakeholders:

He [Michie] had this tiny grant from the University for his research unit [...] even though he was sidelined, he was perfectly capable of raising funds from industrial outfits like BP. So, you know, he was always perfectly okay. And even SRC gave him some grants. So, he was fine. And ridiculously self-confident. (Blake, 2024).

At its formation in 1974 DAI had six academic staff: Bernard Meltzer, Alan Bundy, Rod Burstall, Jim Howe, Gordon Plotkin and Robin Popplestone (Meltzer, 1975). In 1979 Burstall and Plotkin moved to the Theory Group in the Department of Computer Science.

The years after the Lighthill report were a lean time for AI in the UK. However, in 1982 there were two reports that were supportive of accelerated R&D into AI. Peter Swinnerton-Dyer and Roger Needham (both of Cambridge University) were commissioned by the Science and Engineering Research Council (SERC, the successor to the SRC) in 1982 to report on the state of AI in the UK (information from Oakley & Owen, 1990, pp. 49-50). And the Alvey Report (Alvey, 1982) for the Department of Trade and Industry on "A Programme for Advanced Information Technology" recommended work on Intelligent Knowledge Based Systems as one of the four key enabling technologies identified. For some, this was the UK response to the Japanese Fifth Generation Computer Systems project that had allegedly triggered the American response of the Strategic Computing Initiative (Feigenbaum & McCorduck, 1983; Roland & Shiman, 2002). The Alvey initiative from 1983 onwards allowed a rapid expansion in staffing, particularly in the area of knowledge-based expert systems.

This endorsement of AI research in the UK was largely seen as a “reversal” of the Lighthill report:

Shortly after 1980, maybe 81 or 82, they decided to reverse this AI winter by getting in a figure as big as Lighthill, which was Sir Peter Swinnerton-Dyer. [...] So he was the anti Lighthill. He was the antidote, I should say, to Lighthill. So Lighthill started the AI winter, then they commissioned Sir Peter Swinnerton-Dyer to reverse it. And why did they do that? Because the Japanese had started something that they called the Fourth [sic] Generation of Computing. (Blake, 2024).

This is further verified according to the director of the programme Brian Oakley, alluding to poet John Milton, “[if] the Lighthill Report of the early 1970s was paradise lost for the AI community, the Alvey Report of the early 1980s was paradise regained” (Oakley, 1990, pp. 347-348). Nevertheless, the peculiar effects of the AI winter were still felt. As an AI expert who took part in the Alvey programme stated in a recent interview:

I’ve never seen so much research funding in my life again. It was kind of curious. [...] I moved into this project and this project was essentially about trying to use AI techniques, knowledge-based system techniques to support designers. We call this an intelligent knowledge based design support system. Because back in those days in the Alvey programme we weren’t allowed to call it AI because this was still post-AI winter stuff (anonymized respondent, cited in Galanos, 2023, p. 117).

Another crucial point for that era was that while neural net research has already been conducted at Edinburgh (in particular, work on associative memory nets, or ‘Willshaw nets’ by Willshaw et al., 1969), the same period marked the publication of Minsky and Papert’s book *Perceptrons*, a scathing critique of limitations the neural network, or connectionist, approach to AI (Minsky & Papert, 1969). This has proved important for research directions at Edinburgh. For example, according to the statement by Geoffrey Hinton who embarked on his doctoral journey in 1972 at Edinburgh:

I went to Edinburgh, and I studied with somebody called Christopher Longuet-Higgins, who was a very clever guy, and had worked on neural networks. But, just before I got there to work with him, he decided neural networks were nonsense. So I spent my five years arguing with him. He kept trying to persuade me to stop doing neural networks and do more conventional symbolic AI. [...] the *Perceptrons* book came out in ’69 and that pretty much destroyed the field.’ (Hinton in interview with Ashlee Vance, Vance, 2021).

In a more recent interview, Hinton ascribes Longuet-Higgins’s steer towards symbolic AI as due to him being convinced by Winograd’s PhD work (Winograd, 1971) on a system that allowed a user to interact with a (virtual) blocks world via a conversation in English, and suggests that social ties and conventions were the only reason that enabled him to complete his PhD at Edinburgh:

At about the same time as I arrived, he gave up on neural networks. He was convinced by Winograd’s thesis, among other things, that neural networks were the wrong approach and symbolic AI was the right approach. He never liked probabilistic things. He liked things to be neat. He was a mathematician and liked clean, neat things. He didn’t like probabilities. And so he and I didn’t get along very well. I now think retrospectively, the only reason I survived as a graduate student was because my father was a Fellow of the Royal Society. So he was in the same club as Longuet-Higgins. And it would have been embarrassing for him to get rid of someone whose father was in the Royal Society. (Hinton, 2024).

In 1985 the School of Epistemics became the Centre for Cognitive Science within the Faculty of Science, devoted exclusively to research and postgraduate teaching. This became the base for Edinburgh’s future strength in natural language processing. In 1998, the University joined together the Departments of Artificial Intelligence and Computer Science, the Centre for Cognitive Science and a number of research institutes to form the Division of Informatics, which later became the School of Informatics.

3 Research Highlights

This section briefly describes early work on reinforcement learning, search, robotics, programming, automated reasoning, planning, natural language processing, computer-based learning environments and cognitive science carried out in Edinburgh. The selection of material is based largely on Alan Bundy’s talk (Bundy, 2023), enhanced by conversations with Alan Bundy and Mark Steedman. It is important to note that the presence of all of these areas in one centre and the strong interdisciplinary culture (outlined above) provided a critical mass and a rich environment for cross-fertilization.

Reinforcement Learning: The Matchbox Educable Noughts And Crosses Engine (MENACE) (Michie, 1961, 1963) was a very early example of reinforcement learning. It played a game of noughts-and-crosses (or tic-tac-toe). It took the first turn, and then alternated turns against a human player. Depending on whether the machine won, drew or lost a game, it obtained positive or negative reinforcement for the sequence of moves it made. MENACE was able to successfully learn how to perform well on the game. It is interesting that the learning algorithm was actually implemented at first using a matchbox for each state, due to a lack of access to a digital computer.

Michie and Chambers (1968) further developed reinforcement learning to tackle the problem of learning to balance a pole hinged to a movable cart. Sutton and Barto (1998, sec. 1.6) describe this as “one of the best early examples of a reinforcement learning task under conditions of incomplete knowledge. It influenced much later work in reinforcement learning, beginning with some of our own studies”.

Search: Many problems in AI can be addressed by search, e.g., puzzles and game playing. However, the search space can grow rapidly with the problem size, so that optimal solutions are very expensive to compute. This leads to the use of heuristics which aim to obtain good solutions for a small fraction of the cost compared to the optimal solution. Graph Traverser (Doran & Michie, 1966) was an early work on heuristic search. As Nilsson (2010, sec. 12.1.1) details, this work helped to inspire a group at the Stanford Research Institute to develop the celebrated A* search algorithm (Hart et al., 1968). Their key contributions were to combine a heuristic estimate of the cost to the goal with the current path length, and to prove that A* is guaranteed to find an optimal path when provided with a consistent heuristic.

Robotics: Freddy (1969–1971) and Freddy II (1973–1976) were experimental robots built at 5 Forrest Hill in Edinburgh. Freddy II was a “hand–eye” robot that could assemble toy wooden models from a heap of pieces. It used vision to identify and locate the parts, and was able to re-arrange them to enable identification when they were obscured by other parts (Ambler et al., 1975). Given the state-of-the-art at the time, this required not only building the robot, but also designing and building the vision system, and a programming environment for controlling the various subsystems. The team involved with Freddy included Patricia (“Pat”) Ambler, Yin Baolin, Harry Barrow, Bob Beattie, Ilona Bellos, Chris Brown, Rod Burstall, Steve Cameron, Gregan Crawford, Roy Featherstone, Anastasia Koutsou, Donald Michie, Robin Popplestone, Stephen Salter,³ and Ken Turner.⁴

Later work by members of the team developed RAPT (Popplestone et al., 1978), providing a higher-level specification for robot behaviour at the object level, rather than low-level actuator programming. Freddy II has been on display at the National Museum of Scotland since 2006.

Programming: The programming language POP-2 (Burstall et al., 1968) was designed and developed by Robin Popplestone and Rod Burstall. It supported much subsequent UK research and teaching in AI. Later, Warren et al. (1977) developed a Prolog compiler for the DECsystem-10 (written in Prolog). This robust implementation paved the way for the widespread use of Prolog.

Automated Reasoning: Meltzer assembled a stellar group of researchers in automated reasoning, including Robert Boyer, Alan Bundy, Patrick Hayes, Robert Kowalski, J. Strother Moore and Gordon Plotkin. Kowalski and Kuehner (1971) studied “Linear resolution with selection function” (aka SL-resolution) as an inference system for first-order logic. A restricted version of SL-resolution forms the basis for Prolog.

Plotkin (1971) worked on inductive inference (or learning from examples), generating hypotheses which are generalisations of experiences. He introduced the important notion of a *least general generalisation* in the framework of first order logic. Boyer and Moore (1973)’s work on proving theorems about LISP functions was notable for its automation of mathematical induction.

Bundy et al. (1979) developed the largest Prolog program of the time called MECHO to solve high-school level mechanics problems specified in predicate calculus and English. Notably MECHO used meta-level inference to control search in natural language understanding, common sense inference, model formation and algebraic manipulation.

Planning: Austin Tate worked on planning for his PhD (supervised by Donald Michie) and subsequently on the Nonlin hierarchical partial-order AI planning system. Nonlin took a joint AI and Operations Research approach to planning, with Lesley Daniel being the OR researcher (Tate, 1977; Tate & Daniel, 1982).

Natural Language Processing: Thorne et al. (1968) was a ground-breaking development in the syntactic analysis of sentences, using an augmented transition network (ATN). This work was an important precursor to William Woods’ famous ATN parser (Woods, 1970).

Going beyond parsing, language exists to enable communication between agents. Davey (1978) built a program to generate a description of a small model universe, specifically the moves in a game of noughts-and-crosses, which it played with its operator. The key problem addressed was to generate explanatory commentaries, referring to entities and moves in terms of their strategic significance in the game, a precursor to modern work in “explanatory AI”. The work was supervised by Christopher Longuet-Higgins and Stephen Isard.

Power (1979) built a program where two robots hold a conversation in order to accomplish a mutual goal, in a world of few objects. The robots could carry out several types of exchange, such as agreeing plans, or obtaining information. The idea behind having a robot–robot interaction was to avoid a human guiding the structure of the dialogue, as they often do in a human–robot conversation.

Computer-based learning environments: At MIT Papert (1973) had developed learning environment to enable a child to communicate with a device called a “turtle” via a program written in LOGO. This was to investigate his view that a child learns by actively exploring their environment. In Edinburgh Jim Howe set up a laboratory to investigate these ideas, working with, *inter alia*, Ben du Boulay, Tim O’Shea and Sylvia Weir. See Howe and O’Shea (1978) for an evaluation of the effects of LOGO programming on learning a number of mathematical topics, and Emanuel and Weir (1976) for a study of how controlling the LOGO turtle helped the development of language for communication in an autistic child.

Cognitive Science: Longuet-Higgins’ group worked on topics across a wide range of areas in addition to the work on natural language of Davey (1978) and Power (1979). The associative net of Willshaw et al. (1969) was an early neural network model which learned to associate input–output pairs of patterns with a Hebbian learning rule. Longuet-Higgins was a fine musician, and worked on topics such as the computational representation and analysis of harmony and metre in music (see, e.g., Longuet-Higgins & Steedman, 1971). Buneman (1971) tackled the problem of reconstructing trees from dissimilarity data, which could be applied to tracing copying errors in manuscripts, as well as to reconstructing phylogenetic trees from genetic data. Longuet-Higgins also had interests in vision, and e.g. supervised Geoffrey Hinton’s PhD thesis (Hinton, 1977) on how the best consistent combination from among many parts or aspects of visual input may be obtained by constraint relaxation.

4 Women in AI Research

It is well documented that certain social demographics have been and are under-represented in STEM subjects (All-Party Parliamentary Group on Diversity and Inclusion in STEM, 2021), for example with respect to gender, ethnicity and economic disadvantage. Here we focus on women in AI, where there were a number of notable female pioneers at Edinburgh. For instance, Pat Ambler joined the EPU as a research assistant in 1966, and developed her career in Edinburgh, becoming a Lecturer in DAI in 1984.⁵ Her work was on robotics, notably Freddy-II (Ambler et al., 1975) and RAPT (Popplestone et al., 1978). She moved to the University of Aberdeen in 1986, and retired from there in 1999 (Sleeman, 2017).

Dr Sylvia Weir was working as an RA with Dr Horace Townsend (clinical neurophysiologist and part-time Senior Lecturer in Artificial Intelligence 1968–1979) from the late 1960s.⁶ She joined DAI in 1974 as a member of Jim Howe’s LOGO laboratory.⁷ Weir studied the catalysis of communication in autistic children in a LOGO-like environment (Emanuel & Weir, 1976), and was also a member of the team that wrote the early AI textbook (Bundy et al., 1978). She was recruited by Seymour Papert to MIT in 1978, and continued to work on using LOGO programming for nurturing of a variety of learning styles, especially for persons with special needs, resulting in the book (Weir, 1987).

In 1985 DAI graduated its first two female PhD students, Martha Palmer (natural language) and Helen Pain (AI in education). Subsequently there have been many notable female PhD students, researchers and academic staff in AI in Edinburgh.

5 Commercialization

In 1969 Michie and Howe established a small company called Conversational Software Ltd (CSL) to develop and market the POP-2 programming language. Howe (2007) relates that the development costs for CSL exceeded the market demand, but that CSL’s POP-2 system supported work in industry and academia for a decade or more, long after the company ceased to trade. There were other commercialization activities over the intervening years, but with the expansion of interest in AI in the early 1980s, DAI set up a separate non-profit organization the Artificial Intelligence Applications Institute (AIAI) in 1983, under the initial directorship of Jim Howe, succeeded by Austin Tate a year later. Much effort was given to developing and testing mechanisms for imparting knowledge and skills to its clients. AIAI’s work focused on three sub-areas of knowledge-based systems: planning and scheduling systems, decision support systems and information systems. Overseas, it had major clients in the US and Japan. A 10-year snapshot of its activities showed it had an annual turnover of just less than £1M, and had broken even from the start. In 2001 AIAI became part of the Centre for Intelligent Systems and their Applications (CISA) within the School of Informatics. Materials about its work can be found at <https://www.ai.ai.ed.ac.uk/index-prev.html>.

A particular case in point for Edinburgh’s unique position in the global AI landscape in the period of AIAI’s formation was that while there was a general (American) sentiment that the Japanese developments posed a “threat” (Feigenbaum & McCorduck, 1983) to AI competitors. Edinburgh’s AIAI, while maintaining friendly contacts with the US expert system

community (and Feigenbaum in particular) also exported the Prolog computing language to Japan for the development of Fifth Generation Computer Systems:

And my contacts were mostly with people like Ed Feigenbaum on that, at Stanford. So we were likewise in Edinburgh doing similar things through AIAI and other groups. So, that definitely gave us the opportunity to start growing. The other influence though was that there was the work on the Fifth Generation Computing Programme in Japan. And they adopted a sort of declarative engine approach which suited logic programming, suited Prolog, and Edinburgh was well placed there. We had a lot of clients in Japan, Edinburgh Prolog was well marketed into Japan and taken up and used by companies and academic groups. And it was Edinburgh Prolog, it wasn't, you know, the variants of Prolog, [...] (Tate, 2024).

This was paired with efforts in tailoring AI research knowledge to commercial partners' ideas (including commercial ideas) while fostering a network of pioneering researchers with Edinburgh as its landmark point of reference:

So it's very typical of what we were doing. We taught hundreds of people, trained hundreds of people, let them do their own projects, but alongside our staff. [...] Had to get them quickly into a position where they could do an application that was relevant to their business or indeed was the thing the business wanted to build, and prototype it alongside our staff, with our staff giving them the systems and the help that they would get if they were doing an MSc. That was the model we had. And we had hundreds of people through that programme. So it was quite effective to make sure all of that could be done. (Tate, 2024).

Such approaches to teaching as part of the commercialization efforts were thus complementing the multiple connections fostered between Edinburgh, the UK, and internationally. A particularly important aspect of Edinburgh's AI research culture was the establishment of the AI departmental library, initiated by Michie's vision to concentrate relevant materials for the various groups' perusal. It became a pole of attraction for international researchers and collaborators, providing "early access to papers that were technical reports as well as published papers" (Tate, 2024), and served both local and external users through a professional library service. Tragically, the library and the newly acquired Turing Institute collection were lost to a fire in 2002, erasing much of this unique resource.

6 Teaching

As per Howe (2007)'s description of teaching programmes at Edinburgh, there had been Masters and PhD level teaching during the 1960s; for example a Diploma course in Machine Intelligence Studies (with 5 students) was first delivered in 1965 (Alonso, 2024). As part of the formation of DAI in 1974, the Department committed itself to contributing to undergraduate teaching. The first undergraduate course offered was AI2 (a second year course) in 1974/75. The material was published in a book (Bundy et al., 1978), one of the first undergraduate AI courses in the world. The book covered the topics of knowledge representation, natural language, visual perception, learning, and programming (in LOGO). Subsequently an introductory course, AI1, was introduced in 1978/79.

The 4-year undergraduate degree in Edinburgh was (and is) made up of two pre-honours years (years 1 and 2), and two honours years (years 3 and 4). In 1982 DAI launched its first joint degree Linguistics with Artificial Intelligence, Artificial Intelligence and Computer Science was started in 1987, and Artificial Intelligence and Mathematics in 1992. A number of other joint degrees were developed, as well as a single-honours degree in Artificial Intelligence.

A Masters degree in Knowledge Based Systems was established in 1983. This offered specialisms in the Foundations of AI, Expert Systems, Intelligent Robotics and Natural Language Processing. There were 40 to 50 graduates produced annually from this course, with many of the Department's PhD students coming from its ranks.

According to researchers who have attended these courses, the breadth of coverage was significant for their career, as it enabled them to be familiar with nearly all aspects of the broad AI research spectrum. For example, as Hallam recollects from preparatory courses for the PhD program during the period 1979–80:

It was fairly informal. There was a set of general kind of introductory lectures on AI. Because there weren't that many textbooks at that time, and basically, no one knew what it was, except at AI departments. So I think we had one or two mornings a week we spent just being brought up to speed on the different subfields of the discipline. And that was very interesting, because that covered the whole breadth of AI. [...] I think that the thing I appreciated most I think about those years was the breadth of coverage that we had. And one of the things I've noticed all through my career is that the breadth of AI that I'm comfortable in is considerably wider than many people's. And that goes back to us having a solid grounding in each of the different sub-disciplines. (Hallam, 2024).

7 National and International Connections

From the start, there were strong connections between AI at Edinburgh and groups in the USA at MIT, Stanford and CMU. Both Michie and Meltzer had visited AI projects there in the early 1960s. In 1965 Michie organised a Machine Intelligence Workshop; this was the first in a series of seven such workshops, held annually in Edinburgh, which attracted leading AI researchers from the United States as well as the UK. The workshop series continued on after this at venues around the world, including the USA, USSR and Japan. The workshops were of great importance for the social development of the field, and helped to establish Edinburgh as an AI centre with an international reputation. There was considerable exchange of personnel between the various centres. Proceedings of the workshops were published, and provided an accessible introduction to the field.⁸

For people like Gordon Plotkin, Austin Tate, and Andrew Blake, this series was pivotal for their choice of Edinburgh as a research base or for piquing their general interest in AI. For Austin Tate, these volumes acquired Christmas gift value status:

When I was at secondary school, I remember getting my parents to buy me the volumes of the machine intelligence workshops. [...] They're lovely volumes that produce really a good summary of AI at the time, with all the different kinds of papers in it. And I had the first volumes of that having got my parents to buy me that for one Christmas present. (Tate, 2024).

Edinburgh AI connections expanded across the UK. For example, the British Computer Society Study Group on Artificial Intelligence and the Simulation of Behaviour (AISB) was founded in 1964. The initial steering committee consisted of Donald Broadbent (Cambridge), Irving John Good (Oxford), Donald Michie (Edinburgh), Christopher Strachey (Oxford) and Max Clowes (Oxford). By May 1965 there were 60 members. In 1974 the study group transformed into the Society for Artificial Intelligence and the Simulation of Behaviour, independent of the BCS. The main activities of the AISB were to arrange scientific meetings and to publish a newsletter. Edinburgh researchers were heavily involved: in addition to Michie's role, Meltzer was Chair of AISB 1972-1976, and Burstall, Doran, Hayes, Bundy and O'Shea edited the newsletter at various times. The AISB was not exclusively British, and indeed following a meeting of European delegates at the first IJCAI in 1969, the newsletter was entitled the European Newsletter for AISB. For a while the society functioned "as a kind of European AISB, until other national groups [were] formed on the continent". Eventually in 1982 the European Coordinating Committee for Artificial Intelligence (ECCAI, now EurAI) was founded with the AISB as a member.⁹

As Fleck (1987) has documented, AI developed at other centres in the UK in the early 1970s, notably at Essex and Sussex, with some influence from Edinburgh. At Essex, J. M. (Mike) Brady was appointed a Lecturer in 1970, and developed interests in image analysis. He was joined by Hayes (1972) and Doran (1973) from Edinburgh, along with other AI researchers. At Sussex, N. S. (Stuart) Sutherland created a Cognitive Studies Programme in the early 1970s. In 1973 Max Clowes took up the Chair of AI there, and in 1974 Longuet-Higgins and some members of his research group moved to Sussex from Edinburgh. Other notable researchers at Sussex at the time included Margaret Boden, Philip Johnson-Laird and Chris Darwin. While in the 1970s there were only a few centres of AI research in the UK, by the mid-1980s (post-Alvey) many higher education institutions in the UK had a serious interest in AI research and training.

8 Discussions on Ethical and Responsible Uses of AI

The contemporary establishment of responsible AI and AI ethics approaches (see, e.g., Dignum, 2023) encourages us to look at early instances of relevant traits in Edinburgh AI circles. Richard Gregory gave a deliberately provocative talk on "The Social Implications of Intelligent Machines" at an Edinburgh-based Machine Intelligence Workshop (Gregory, 1971, reprinted in Gregory, 1974). In that speech he touches on the large unpredictability of future states of entities with intelligence (given that an aspect of intelligence is the escape from predictable states); issues of accountability and agency in decision-making machines (also speculating on robotic judges that have become more fashionable in the last decade); questions concerning human-machine relations; issues of privacy arising from the large-scale storage of personal data; difficulties in instilling moral values in machines especially given the observation that morals change across periods and regions in human societies; and concluding that a world with intelligent machines could be compatible with human happiness if the machines are very different from us, or show servitude of the machine to humans, akin to that of the slave to the master. A similar approach, and with adjacent vocabulary can be found at Michie's 1968 paper "Computer – servant or master?" in which Michie outlines the current state of AI machines and how their "intelligence" and "learning" differs to humans', while making a case for cooperation of assistance on behalf of the machine toward the human (Michie, 1986, pp. 240-246).

Michie advanced the discussion of ethical issues in AI by organizing, in 1972, a small meeting of ten persons at the Villa Serbelloni, Lake Como, Italy on "The Social Implications of Machine Intelligence Research". Participants included,

amongst others, the notable US AI scientists John McCarthy and Daniel Bobrow. Powell (2023) provides an in-depth discussion of the meeting and its outcomes, and notes that while the group identified a “a wide range of increasingly important social consequences”, they did not advocate for action, believing that more time was needed for AI to be “developed to any significant degree”. This connects to Collingridge’s view about the “dilemma of control” in the regulation of emerging technologies (Collingridge, 1982); too early regulation may hamper beneficial innovation, but regulation may come too late if a lot of innovation paired with infrastructural entrenchment is embedded across numerous social strata.

There were early concerns about military funding for research. For example a letter from Peter Buneman and Stephen Isard to the *University Bulletin* (Buneman & Isard, 1972) expressed unease about the dominance of funding from the US Department of Defense Agency ARPA for American Research in AI. They state that “Twenty-five of the School [of Artificial Intelligence]’s members have recently signed a document in which they declared their opposition to any financial dealings between the School and ARPA. We feel that there is a danger that the whole field of Artificial Intelligence may be distorted towards military ends [...]”. Sloman (2011) recollects conversations about “the ethics of using AI in development of weapons” at a 1979 Edinburgh-based Machine Intelligence Workshop. For John Hallam, the conversations about ethics were present, but they were paired with the difficulty of predicting future developments in AI technology. He reports that the US Strategic Defense Initiative (announced in 1983) was very poorly received, but that general computing funding carried a sense of technological neutrality.

Prompt: During your stay in Edinburgh, did you have the sort of conversations around political influences or misuses of AI? Could you envision where AI and robotics might end up in two or three decades?

Hallam: I think no. I don’t mean no, we didn’t have the conversation, I think no, we realised it was practically impossible to envision where what we were doing might end up. [...] But we did have conversations about what AI might not do. And everybody got very cross with Reagan and his Star Wars initiative and refused to work on it. Pretty much the whole of the AI community said no, we’re having nothing to do with this. [...] It was the Strategic Defense Initiative, which was the Star Wars programme, that particularly worried people. I think having lots of money to develop interesting computing technology, that wasn’t an issue. Because people viewed technology, I think much more neutrally. People worry much more now about what technology will do. But it’s a very difficult question. (Hallam, 2024).

Indeed there was quite a lot of activity of various people in DAI in the early 1980s in response to the Strategic Defense Initiative. For example Henry Thompson and Computing & Social Responsibility (1985) wrote the paper “There Will Always Be Another Moonrise” about the dangers of relying on automated decision making in such contexts. The title is based on the false alarm in 1960 created by an early warning system in Greenland that a missile attack had been launched against the United States. It turned out that radars, apparently, had echoed off the moon.

The current range of debates about Ethics and AI can be summarised from the headings in the article by Müller (2023) in the Stanford Encyclopedia of Philosophy. This lists ten areas: (i) Privacy & Surveillance, (ii) Manipulation of Behaviour, (iii) Opacity of AI Systems, (iv) Bias in Decision Systems, (v) Human–Robot Interaction, (vi) Automation and Employment, (vii) Autonomous Systems, (viii) Machine Ethics, (ix) Artificial Moral Agents, and (x) Singularity. Of these we can detect at least some coverage of areas (i), (iii), (v), (vi), (vii) and (ix) from the early years of AI work in Edinburgh.

9 Conclusions

This article has presented a comprehensive exploration of the foundations of AI at the University of Edinburgh from the 1960s to the mid-1980s. We have charted the establishment of research groups and departments, highlighting the significant contributions and charismatic influence of pivotal figures such as Donald Michie, Christopher Longuet-Higgins, Richard Gregory and Bernard Meltzer among others. We have highlighted notable female pioneers of AI in Edinburgh. Our narrative has captured the interdisciplinary nature of the AI initiatives at Edinburgh, supported by ongoing national and international collaborations that positioned Edinburgh as a European AI hub. The analysis has incorporated both historical accounts and recent oral histories, facilitating a nuanced understanding of events and their legacy. Our analysis has revealed the challenges faced, including the controversial “AI winter” of the early 1970s. This period was characterised by scepticism towards AI’s industrial applicability (as per the 1973 Lighthill report), rather than merely its technical shortfalls. Our findings suggest that social dynamics, including scientific egoism and rivalry, contributed to the challenges faced. Yet AI research at Edinburgh persisted, and further diversified in response to evolving funding landscapes, such as through the Alvey initiative in the early 1980s.

Technical accomplishments from this period underscore the era’s innovative spirit. Notable too was the establishment of the AI departmental library, catalysing knowledge exchange despite its loss to fire in later years. On the teaching front, the

breadth of AI courses offered at Edinburgh in the 1970s and 1980s underscored the formative role it played in developing knowledgeable and versatile AI practitioners. Further, this article has underscored the early ethical discourse around AI, set against a backdrop of views about technological neutrality. It demonstrates some early reflections on AI's ethical use, albeit constrained by the era's limited foresight into AI's future societal impact.

Our historical account is inevitably incomplete. More historical research is to be done on a number of areas, including:

1. How did Edinburgh influence other areas in Europe and beyond, in addition to what we have recorded?
2. What other elements, such as global historical or economic factors, could contribute to a better assessment of the flourishing of AI at Edinburgh, as well as the "AI winter" and the Lighthill controversy?
3. Focus should not be restricted to the "big four" founding Professors; there were many other figures in Edinburgh who made major contributions to the development of the field, and who are worthy of further study.
4. What were the relations between the Department of Computer Science (founded in 1966) and the various AI departments?
5. How does the early history that we have covered (ca. 1960-1989) connect to more recent developments at Edinburgh and beyond?

Overall, this historical examination not only clarifies Edinburgh's foundational role in AI's evolution but also exposes enduring themes critical to modern AI discourse, especially around ethical responsibility and the inter-disciplinary, borderless nature of AI science.

More resources about the history of AI in Edinburgh can be found at https://groups.inf.ed.ac.uk/aics_history/.

Acknowledgements

Many thanks to Alan Bundy, Peter Buneman, Bob Fisher, Greg Michaelson, Gordon Plotkin, Mark Steedman, Austin Tate, Henry Thompson and David Willshaw for their input, and to the anonymous referees whose comments helped to improve this article.

For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

ORCID iDs

Christopher KI Williams  <https://orcid.org/0000-0002-6270-4703>

Vassilis Galanos  <https://orcid.org/0000-0002-8363-4855>

Xiao Yang  <https://orcid.org/0000-0003-4389-6495>

Funding

The authors thank the School of Informatics for providing support part-time for Xiao Yang Feb-July 2024 for transcription of interviews and creation of the website.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Notes

1. Gregory's CV is available at <https://www.richardgregory.org/cv.htm>
2. *University Bulletin* 3(4) Jan 1967 gives the date as 1 Oct 1966. See the obituary at <https://www.inf.ed.ac.uk/events/news/obituary/bmeltzer.html> for more information about Meltzer.
3. Later famous for his work on wave power.
4. List from <https://groups.inf.ed.ac.uk/vision/DATASETS/ROBOTICS/FREDDY/ambler1.htm>
5. *University Bulletin* 21(3) 1984.
6. Pers. comm. Prof D Willshaw, 19/2/25.
7. Sylvia Weir obituary: <https://www.theguardian.com/education/2018/oct/29/sylvia-weir-obituary>.
8. Details of the Machine Intelligence book series are available at <https://www.doc.ic.ac.uk/~shm/MI/mi.html>.
9. This paragraph is based on information in Alonso (2024).

References

- Agar, J. (2020). What is science for? The Lighthill report on artificial intelligence reinterpreted. *British Journal for the History of Science*, 53(3), 289–310.
- All-Party Parliamentary Group on Diversity and Inclusion in STEM. (2021) Inquiry on equity in the STEM workforce: Final report.

- Alonso, E. (2024). The early days of the AISB. In U. Cortés, A. Bugarín, & C. Barrué (Eds.), *Proceeding of the Workshop on the History of AI in Europe (WHAI@EU)* (pp. 3–6). Zenodo.
- Alvey, J. (1982). *A programme for advanced information technology*. https://www.chilton-computing.org.uk/inf/literature/reports/alvey_report/overview.htm
- Ambler, A. P., Barrow, H. G., Brown, C. M., Burstall, R. M., & Popplestone, R. J. (1975). A versatile system for computer-controlled assembly. *Artificial Intelligence*, 6, 129–156.
- Birse, R. M. (1994). *Science at the University of Edinburgh 1583–1993*. The University of Edinburgh.
- Blake, A. (2024). Interview. Interviewer: V. Galanos. https://groups.inf.ed.ac.uk/aics_history/oral_history.html
- Boyer, R. S., & Moore, J. S. (1973). Proving theorems about LISP functions. In *Proceedings of the third international joint conference on artificial intelligence (IJCAI-73)* (pp. 486–493).
- Bundy, A. (2023). The early years of AI & cognitive science in Edinburgh. https://media.ed.ac.uk/media/1_78nvtbyg
- Bundy, A., Burstall, R. M., Weir, S., & Young, R. M. (1978). *Artificial Intelligence: An introductory course*. Edinburgh University Press.
- Bundy, A., Byrd, L., Luger, G., Mellish, C., & Palmer, M. (1979). Solving mechanics problems using meta-level inference. In *Proceedings of the sixth international joint conference on artificial intelligence (IJCAI-79)* (Vol. 2, pp. 1017–1027).
- Buneman, O. P., & Isard, S. D. (1972). “Social uses” of machine intelligence. *University of Edinburgh Bulletin*, 8(14), 8.
- Buneman, P. (1971). The recovery of trees from measures of dissimilarity. In F. R. Hodson, D. Kendall, & P. Tăutu (Eds.), *Mathematics in the archaeological and historical sciences* (pp. 387–395). Edinburgh University Press.
- Burstall, R., Collins, J., & Popplestone, R. (1968). *The POP-2 papers*. Edinburgh University Press.
- Collingridge, D. (1982). *The social control of technology*. St Martin’s Press.
- Copeland, B. J. (2023). Early AI in Britain: Turing et al. *IEEE Annals of the History of Computing*, 45(3), 19–31.
- Crevier, D. (1993). *AI: The tumultuous history of the search for artificial intelligence*. Basic Books.
- Davey, A. (1978). *Discourse production*. Edinburgh University Press.
- Dignum, V. (2023). Responsible artificial intelligence: Recommendations and lessons learned. In D. O. Eke, K. Wakunuma, & S. Akintoye (Eds.), *Responsible AI in Africa: Challenges and opportunities* (pp. 195–214). Springer International Publishing.
- Doran, J. E., & Michie, D. (1966). Experiments with the graph traverser program. *Proceedings of the Royal Society, A*, 294, 235–259.
- Emanuel, R., & Weir, S. (1976). Catalysing communication in an autistic child in a LOGO-like learning environment. In *Proceedings of the 2nd Summer conference on artificial intelligence and simulation of behaviour* (pp. 118–129).
- Feigenbaum, E. A., & McCorduck, P. (1983). *The fifth generation: Artificial intelligence and Japan’s computer challenge to the world*. Pan Books.
- Fleck, J. (1987). Development and establishment in artificial intelligence. In B. Bloomfield (Ed.) *The question of AI* (pp. 106–164). Croom Helm.
- Galanos, V. (2023). *Expectations and expertise in artificial intelligence: Specialist views and historical perspectives on conceptualisation, promise, and funding*. Ph. D. Thesis, University of Edinburgh.
- Gregory, R. L. (1971). The social implications of intelligent machines. In B. Meltzer & D. Michie (Eds.), *Machine intelligence 6* (pp. 3–13). Edinburgh University Press.
- Gregory, R. L. (1974). *Concepts and mechanisms of perception*. Charles Scribner’s Sons.
- Gregory, R. L., & Murrell, J. N. (2006). Hugh Christopher Longuet-Higgins. *Biographical Memoirs of Fellows of the Royal Society*, 52, 149–166.
- Haigh, T. (2023). There was no ‘first AI winter’. *Communications of the ACM*, 66(12), 35–39.
- Hallam, J. (2024). Interview. Interviewer: V. Galanos. https://groups.inf.ed.ac.uk/aics_history/oral_history.html
- Hart, P. E., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2), 100–107.
- Hinton, G. E. (1977). *Relaxation and its role in vision*. Ph. D. Thesis, University of Edinburgh.
- Hinton, G. E. (2024). Interview. Interviewer: V. Galanos. https://groups.inf.ed.ac.uk/aics_history/oral_history.html
- Howe, J. (2007). *Artificial intelligence at Edinburgh University: A perspective*. <https://www.inf.ed.ac.uk/about/AIhistory.html>
- Howe, J., & O’Shea, T. (1978). Learning mathematics through LOGO. *ACM SIGCUE Bulletin*, 12(1), 2–11.
- Kowalski, R., & Kuehner, D. (1971). Linear resolution with selection function. *Artificial Intelligence*, 2, 227–260.
- Lee, J. (2012). *Donald Michie*. <https://history.computer.org/pioneers/michie.html>
- Lighthill, J. (1973). *Artificial intelligence: A general survey*. Science Research Council.
- Lincoln, Y. S., & Guba, E. G. (1986). But is it rigorous? Trustworthiness and authenticity in naturalistic evaluation. *New Directions for Program Evaluation*, 1986, 73–84.
- Longuet-Higgins, H. C. (1973). Comments on the Lighthill report and the Sutherland reply. In B. H. Flowers (Ed.), *Lighthill report: Artificial intelligence: A paper symposium*. Science Research Council.
- Longuet-Higgins, H. C. (1977). Epistemics. In A. Bullock & O. Stallybrass (Eds.), *Fontana dictionary of modern thought* (p. 209). Fontana.

- Longuet-Higgins, H. C. (1987). *Mental processes: Studies in cognitive science*. MIT Press.
- Longuet-Higgins, H. C., & Steedman, M. J. (1971). On interpreting Bach. In *Machine intelligence 6* (pp. 221–241). Edinburgh University Press.
- Meltzer, B. (1975). Department of A.I. – University of Edinburgh. *SIGART Newsletter*, (50), 5.
- Michie, D. (1961). Trial and error. In S. A. Barnett & A. McLaren (Eds.), *Science survey, 1961* (pp. 129–145). Penguin. Part 2.
- Michie, D. (1963). Experiments on the mechanisation of game-learning part 1. Characterisation of the model and its parameters. *The Computer Journal*, 6, 232–236.
- Michie, D. (1986). *On machine intelligence (2nd ed.)*. Ellis Horwood Limited.
- Michie, D., & Chambers, R. A. (1968). BOXES: An experiment in adaptive control. In E. Dale & D. Michie (Eds.), *Machine intelligence* (Vol. 2, pp. 137–152). Edinburgh University Press.
- Minsky, M. L., & Papert, S. (1969). *Perceptrons*. MIT Press.
- Moore, J. S. (2021). Recollections of Hope Park Square, 1970–1973. In G. Michaelson (Ed.), *Mathematical reasoning: The history and impact of the DReaM group* (pp. 37–49). Springer.
- Müller, V. C. (2023). Ethics of artificial intelligence and robotics. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*, Fall 2023 ed. Metaphysics Research Lab, Stanford University.
- Nilsson, N. (2010). *The quest for artificial intelligence: A history of ideas and achievements*. Cambridge University Press.
- Oakley, B., & Owen, K. (1990). *Alvey: Britain's strategic computing initiative*. MIT Press.
- Oakley, B. W. (1990). Intelligent knowledge-based systems – AI in the UK. In R. Kurzweil (Ed.), *The age of intelligent machines* (pp. 346–349). MIT Press.
- Olazaran, M. (1996). A sociological study of the official history of the perceptrons controversy. *Social Studies of Science*, 26(3), 611–659.
- Papert, S. (1973). *Uses of technology to enhance education*. MIT AI Memo 298.
- Plotkin, G. D. (1971). *Automatic methods of inductive inference*. Ph. D. Thesis, University of Edinburgh.
- Popplestone, R., Ambler, A., & Bellos, I. (1978). RAPT: A language for describing assemblies. *Industrial Robot*, 5(3), 131–137.
- Powell, R. (2023). The ‘artificial intelligensia’ and its discontents: An exploration of 1970s attitudes to the ‘social responsibility of the machine intelligence worker’. *BJHS Themes*, 8, 189–203.
- Power, R. (1979). The organisation of purposeful dialogues. *Linguistics*, 17, 107–152.
- Roland, A., & Shiman, P. (2002). *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983–1993*. MIT Press.
- Sleeman, D. (2017). Pat Fothergill (Pat Ambler): Early worker in robotics and the use of AI in engineering design. *AISB Quarterly*, (146), 4–8.
- Sloman, A. (2011). John McCarthy – Some reminiscences. <https://www.cs.bham.ac.uk/research/projects/cogaff/sloman-jmc-aisb.pdf>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning*. MIT Press.
- Tate, A. (1977). Generating project networks. In *Proceedings of the fifth international joint conference on artificial intelligence (IJCAI-77)* (pp. 888–893).
- Tate, A. (2024). Interview. Interviewer: V. Galanos. https://groups.inf.ed.ac.uk/aics_history/oral_history.html
- Tate, A., & Daniel, L. (1982). *A retrospective on the Planning: A joint AI/OR Approach Project*. Department of Artificial Intelligence Working Paper 125, Edinburgh.
- Thompson H. & Computing & Social Responsibility (1985) There will always be another moonrise – Computer technology and nuclear weapons. *AISB Quarterly*, 53/54, 21–23.
- Thorne, J., Bratley, P., & Dewar, H. (1968). The syntactic analysis of English by machine. In D. Michie (Ed.), *Machine intelligence* (Vol. 3, pp. 281–309). Edinburgh University Press.
- Timmermans, S., & Tavory, I. (2012). Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), 167–186.
- Vance, A. (2021). Oral history: Geoff Hinton on how AI came to be and what we’re supposed to do with it. <https://www.corememory.com/p/oralhistorywithgeoffhinton>
- Warren, D. H. D., Pereira, L. M., & Pereira, F. (1977). Prolog – The language and its implementation compared with Lisp. *ACM SIGART Bulletin*, 64, 109–115.
- Weir, S. (1987). *Cultivating Minds: A Logo Casebook*. Harper and Row.
- Williams, C. K. I., Galanos, V., & Yang, X. (2024). A short history of the early years of artificial intelligence at Edinburgh. In U. Cortés, A. Bugarín, & C. Barrué (Eds.), *Proceeding of the Workshop on the History of AI in Europe (WHAI@EU)* (pp. 22–23). Zenodo.
- Willshaw, D. J. (2024). Interview. Interviewer: V. Galanos. https://groups.inf.ed.ac.uk/aics_history/oral_history.html
- Willshaw, D. J., Buneman, O., & Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, 222, 960–962.
- Winograd, T. (1971). *Procedures as a representation for data in a computer program for understanding natural language*. Ph. D. Thesis, MIT. Published as AITR-235.
- Woods, W. A. (1970). Transition network grammars for natural language analysis. *Communications of the ACM*, 13(10), 591–606.