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

Christopher K. I. Williams[a,*], **Vassilis Galanos**[b] **and Xiao Yang**[a]

[a]School of Informatics, University of Edinburgh
[b]Lecturer in Digital Work, University of Stirling
ORCID (Christopher K. I. Williams): https://orcid.org/0000-0002-6270-4703, ORCID (Vassilis Galanos): https://orcid.org/0000-0002-8363-4855, ORCID (Xiao Yang): https://orcid.org/0000-0003-4389-6495

This paper provides a brief overview of the the academic structures developed for Artificial Intelligence (AI) research at the University of Edinburgh in the 1960s and 1970s, and some highlights of the ground-breaking work carried out.

1 Structures and Chronology

We give a short overview of the various structures in the University of Edinburgh that were developed to accommodate AI research in Edinburgh, based heavily on Jim Howe's document [10].

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 the Experimental Programming Unit with Michie as Director. 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. The latter two had left Cambridge to come to Edinburgh. 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, and Gregory's the Bionics Research Laboratory. Bernard Meltzer had also set up the Metamathematics Unit in the mid 1960s to pursue research in automated reasoning. At this time 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". When Longuet-Higgins left in 1974, Barry Richards of the Department of Philosophy became the Director of the School.

There were some structural re-organizations in the early 1970s, leading finally 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. In 1983 Michie co-founded the Turing Institute in Glasgow.

Part of DAI's role was to be involved in undergraduate teaching, and the staff developed courses for this. The material was published in a book [4], one of the first undergraduate AI courses in the world.

The years after the Lighthill report (1973) [13] were a lean time for AI in the UK, and by 1979 DAI had only four members of academic staff. However, the Alvey initiative (1983, the UK response to the Japanese 5th Generation Project) allowed a rapid expansion in staffing. 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.

2 Research highlights

This section briefly describes work on reinforcement learning, robotics, programming, automated reasoning, 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 from 2023 [3].

Reinforcement Learning: The Matchbox Educable Noughts And Crosses Engine (MENACE) [15] 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.

Robotics: Freddy (1969–1971) and Freddy II (1973–1976) were experimental robots built 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 rearrange them to enable identification when they were obscured by other parts [1]. 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, Harry Barrow, Chris Brown, Rod Burstall, Gregan

* Corresponding Author. Email: C.K.I.Williams@ed.ac.uk

Crawford, Donald Michie, Robin Popplestone, Stephen Salter[1], and Ken Turner.

Later work by members of the team developed RAPT [17], 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 [6] was designed and developed by Robin Popplestone and Rod Burstall. It supported much subsequent UK research and teaching in AI. Later, Warren et al. [20] 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 & Kuehner's 1971 paper [12] 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. Boyer and Moore [2]'s work on proving theorems about LISP functions was notable for its automation of induction.

Bundy et al. [5] 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.

Natural Language Processing: Thorne et al. [19] was a groundbreaking 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 [22].

Going beyond parsing, language exists to enable communication between agents. Davey [7] 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 [18] 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 [16] 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 [11] for an evaluation of the effects of LOGO programming on learning a number of mathematical topics, and Emanuel and Weir [8] 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. The associative net of Willshaw et al. [21] 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 and Steedman [14]). He also had interests in vision, and e.g. supervised Geoffrey Hinton's PhD thesis [9] on how the best consistent combination from among many parts or aspects of visual input may be obtained by constraint relaxation.

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

[1] Later famous for his work on wave power.

Acknowledgements

Many thanks to Alan Bundy, Mark Steedman, Peter Buneman, Gordon Plotkin, and David Willshaw for their input, and to the School of Informatics for providing support part-time for XY Feb-July 2024.

References

[1] A. P. Ambler, H. G. Barrow, C. M. Brown, R. M. Burstall, and R. J. Popplestone. A Versatile System for Computer-Controlled Assembly. *Artificial Intelligence*, 6:129–156, 1975.

[2] R. S. Boyer and J. S. Moore. Proving theorems about LISP functions. In *Proceedings of the Third International Joint Conference on Artificial Intelligence (IJCAI-73)*, pages 486–493. 1973.

[3] A. Bundy. The Early Years of AI & Cognitive Science in Edinburgh, 2023. `https://media.ed.ac.uk/media/1_78nvtbyg`.

[4] A. Bundy, R. M. Burstall, S. Weir, and R. M. Young. *Artificial Intelligence: An introductory course*. Edinburgh University Press, 1978.

[5] A. Bundy, L. Byrd, G. Luger, C. Mellish, and M. Palmer. Solving Mechanics Problems Using Meta-Level Inference. In *Proceedings of the Sixth International Joint Conference on Artificial Intelligence (IJCAI-79)*, volume 2, pages 1017–1027. 1979.

[6] R. Burstall, J. Collins, and R. Popplestone. *The POP-2 Papers*. Edinburgh University Press, 1968.

[7] A. Davey. *Discourse Production*. Edinburgh University Press, Edinburgh, 1978.

[8] R. Emanuel and S. Weir. 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*, pages 118–129, 1976.

[9] G. E. Hinton. *Relaxation and its Role in Vision*. PhD thesis, University of Edinburgh, 1977.

[10] J. Howe. Artificial Intelligence at Edinburgh University: A Perspective, 2007. https://www.inf.ed.ac.uk/about/AIhistory.html.

[11] J. Howe and T. O'Shea. Learning Mathematics Through LOGO. *ACM SIGCUE Bulletin*, 12(1):2–11, 1978.

[12] R. Kowalski and D. Kuehner. Linear resolution with selection function. *Artificial Intelligence*, 2:227–260, 1971.

[13] J. Lighthill. Artificial intelligence: A general survey, 1973. Science Research Council.

[14] H. C. Longuet-Higgins and M. J. Steedman. On Interpreting Bach. In *Machine Intelligence 6*, pages 221–241. Edinburgh University Press, 1971.

[15] D. Michie. Experiments on the mechanisation of game-learning Part 1. Characterisation of the model and its parameters. *The Computer Journal*, 6:232–236, 1963.

[16] S. Papert. Uses of Technology to Enhance Education, 1973. MIT AI Memo 298.

[17] R. Popplestone, A. Ambler, and I. Bellos. RAPT: A language for describing assemblies. *Industrial Robot*, 5(3):131–137, 1978.

[18] R. Power. The organisation of purposeful dialogues. *Linguistics*, 17: 107–152, 1979.

[19] J. Thorne, P. Bratley, and H. Dewar. The syntactic analysis of English by machine. In D. Michie, editor, *Machine Intelligence*, volume 3, pages 281–309. Edinburgh University Press, Edinburgh, 1968.

[20] D. H. D. Warren, L. M. Pereira, and F. Pereira. Prolog - the language and its implementation compared with Lisp. *ACM SIGART Bulletin*, 64: 109–115, 1977.

[21] D. J. Willshaw, O. Buneman, and H. C. Longuet-Higgins. Non-holographic associative memory. *Nature*, 222:960–962, 1969.

[22] W. A. Woods. Transition network grammars for natural language analysis. *Communications of the ACM*, 13(10):591–606, 1970.