Tokyo-Edinburgh dialogue on robots in artificial intelligence research

H. G. Barrow, D. Michie, R. J. Popplestone, and S. H. Salter

Department of Machine Intelligence and Perception, University of Edinburgh

Integrated systems have become prominent in the research programmes of a number of artificial intelligence laboratories. The Edinburgh machine intelligence group here defines its research and development philosophy in reply to questions put by Japanese robot engineers.

(Received May 1970)

At the Conference of the International Federation of Information Processing Societies, which was held in Edinburgh in 1968, E. A. Feigenbaum of Stanford University, USA, delivered a paper entitled 'Artificial Intelligence: themes in the second decade' (Feigenbaum, 1968). In it he said:

> 'History will record that in 1968, in three major laboratories for AI research, an integrated robot consisted of the following:
> (*a*) a complex receptor (typically a television camera of some sort) sending afferent signals to . . .
> (*b*) a computer of considerable power; a large core memory; a variety of programs for analysing the afferent video signals and making decisions relating to the effectual movement of . . .
> (*c*) a mechanical arm-and-hand manipulator or a motor-driven cart.
>
> The intensive effort being invested in the development of computer controlled hand-eye and eye-cart devices is for me the most unexpected occurrence in AI research in the 1963–8 period.'

Since then research on computer-controlled robots, as a major aid to artificial intelligence research, has proceeded apace, for example in the three laboratories mentioned by Feigenbaum, directed respectively by M. Minsky at MIT, J. McCarthy at Stanford University and C. Rosen at Stanford Research Institute.

Recently, Japanese groups have been entering the field in strength, notably the Electro-technical Laboratory in Tokyo. This laboratory was represented by S. Tsuji on a survey team of robot engineering recently sent on a world tour by the Japan Electronic Industry Association under the leadership of Professor Y. Ukita. The team paid a visit, among other ports of call, to the Department of Machine Intelligence and Perception, University of Edinburgh, and submitted a list of 35 questions concerning the project in progress here. We found it an extremely useful and clarifying exercise to answer these questions, which seem to us wide-ranging and shrewd.

Since the aims and content of artificial intelligence research, and of experimentation with robot devices in particular, are not yet widely known outside a very few specialist groups, there may also be benefit in making the dialogue available to a wider scientific readership. We reproduce the text of the exchange below:

1. General

1. Q) What is the purpose of your research on intelligent robots?

   A) To investigate theoretical principles concerning the design of cognitive systems and to relate these to the theory of programming. To devise adequate methods for the formal description of planning, reasoning, learning and recognition, and for integrating these processes into a functioning whole. In terms of application (long-range) we can envisage a possible use of an intelligent robot as a teaching machine for young children. But our project is a research project, not an application project. Robots for us play the role of test gear for the adequacy of the formal descriptions referred to above.

2. Q) Which do you think most important in your research—scene analysis, problem-solving, dexterous manipulation, voice recognition or something else?

   A) Problem-solving.

3. Q) Do you have a plan for developing any new hardware for manipulators, locomotion machines or special processors for vision?

   A) We plan to use equipment already developed by ourselves and others, and we prefer to simulate locomotion by movement of the robot's world as first suggested to us by Mr. Derek Healy. The present 'world' is a 3 feet diameter sandwich of hardboard and polystyrene which is light and rigid. It rests on three steel balls and is moved by wheels, driven by small stepping motors, mounted on the anchored robot. A pair of bumpers, one in front, one behind, operate two micro-switches to determine contact with obstacles. Our next piece of equipment is a platform 5 feet square which may be moved anywhere in a 10 feet square by flexible drive wires from two servo-motors. The platform can carry weights of 200 lb. and will move at up to 10 inches per second with accelerations of 1/10 g. Various types of hand-eye systems may be hung from a bridge above the platform.

4. Q) We assume that the speed of available digital computers is still too slow for real-time processing of complex artificial intelligence problems. Is this true? If so, do you have any ideas for solving the difficulty?

A) We agree that the speed of available computers is still too slow, especially for sophisticated peripheral processing such as vision. Dedication of satellite processors to sub-tasks (e.g. preprocessing the video signal) is one approach. Special-purpose hardware could of course increase the speed of processing, but it seems doubtful whether it can exhibit behaviour of great logical complexity which a digital computer is capable of doing. An improved instruction set, or more parallel computation (multiprocessor) may yield significant improvements. But the immediate obstacles lie in fundamental problems of software design, rather than in hardware limitations.

5. Q) Which language do you use in robot research, FORTRAN, ALGOL, PL/1, Assembler, LISP or other list processing language? What would be the features of robot-oriented languages?

A) We use POP-2 (Burstall and Collins, 1971; Burstall and Popplestone, 1968). The nearer a programming language is to a fully general mathematical notation, the more open-ended its structure, and the more flexibly adapted to conversational use, then the better the language for robot research. We feel that an ideal robot-oriented language would be one that dealt in relations as well as functions, and would have deductive and inductive capabilities.

6. Q) Can you describe the software hierarchy structure in your robot system?

A) The mechanism of hierarchy is simply that of function call and a typical hierarchy might be (example taken from the vision hierarchy)
top—program for guiding object recognition.
middle—region-finding program and program for matching relational structures.
bottom—eye control program.

7. Q) What performance capability do you predict for intelligent robots in 1975?

A) We expect demonstrations of feasibility before 1975 in the child teaching machine application; that is a system able to recognise and manipulate materials used in teaching children the elements of arithmetic, sets, properties and relations, conservation laws etc.

8. Q) Will there be any chance of applying the newly developed techniques in research on intelligent robots to some industry (for example assembly line) in the near future?

A) We see possible industrial applications in the late 1970s including assembly line. Other conceivable applications are luggage handling at airports, parcel handling and packing, machine tool control and repair, and various exploratory vehicles, e.g. for pipe-laying in deserts, forest clearing in remote areas, ocean-bed work and planetary exploration. Applications for cognitive vehicles will probably remain restricted to work in environments which are essentially intractable.

9. Q) What do you think of the control of many industrial robots by a mini-computer? What level of 'intelligence' would such a computer-robot system have?

A) We would certainly expect to see the control of many 'fixed program' robots by a mini-computer. Such a system would not show much intelligence.

10. Q) May we know the budget and manpower available for your project?

A) We have £500 per annum from the Science Research Council for 'construction of models for on-line control experiments' supplemented by small sums earned as revenue through consultancy and rental of computer time. In addition the GPO Telecommunications Headquarters have awarded a contract for £10,000 over 2 years specifically for the robot research.

The mechanical engineering for our Mark 1 robot, costing about £1,000 to construct, was largely the work of Mr. Steve Salter of the Bionics Research Laboratory of this Department, at that time directed by Professor R. L. Gregory and supported by the Nuffield Foundation. The electronics, interfacing and software have been mainly done in the Experimental Programming Unit by one grant-supported research scientist working part-time on the robot work (Dr. Harry Barrow) and one University Lecturer (Mr. Robin Popplestone). But the work is being carried out in the general context of a large-scale study of machine simulation of learning, cognition and perception, financed on a generous scale by the Science Research Council (£260,000 over 5 years) and by the University of Edinburgh. The POP-2 software and conversational computing system has received support also from the Medical Research Council to the amount of about £70,000 over 5 years. About a dozen research scientists are employed in the general project. Seven of these constitute a 'Robot Working Party' which meets fortnightly under the chairmanship of Professor Donald Michie, and plans the robot work, but this is a side-line activity for them with the exception of the workers mentioned above.

2. Eye

1. Q) What are the aims and targets of your research in the context of vision?

A) Picture-processing performance should be sufficient for forming plausible recognition hypotheses concerning members of a limited repertoire of simple objects (e.g. ball, pencil, cylinder, wedge, doughnut, cup, spectacles, hammer) as a basis for experimental verification or modification of such hypotheses by the robot through action (changing angle of view or interfering with objects manually).

2. Q) Which input device do you use: vidicons, image dissector tubes or other special devices?

A) We use vidicons but are investigating image dissectors.

3. Q) What is the performance of the input devices in areas such as resolution, dynamic range,

sampling rate of A to D converters? In such areas are there any possibilities of improving the input devices?

A) Present resolution of TV sampling system is 64 × 64 points and 16 brightness levels. Speed of conversion of A to D converter is approximately 100 KHz. This system is to be improved to 256 × 256 points and 64, or more, levels. A to D conversion should be about the same rate.

Sampling time for a picture point is largely determined by the time taken for the TV scan to reach the point (up to 20 ms maximum). We are considering image dissectors, which have negligible settling time.

4. Q) Do the eyes of your robot move (electronic or mechanical movement)? What are the merits of eye movement?

A) The eye does not move relative to the main frame. We are considering relative movement of two eyes for depth perception. Also, we are considering using one camera for wide angle views and a second camera with a long-focus lens for investigation of details. Merits—obvious; demerits—complication.

5. Q) Is there any processor for visual input? Is it special hardware? What is the role of the preprocessor?

A) We have installed a small processor for preprocessing visual input and thus reduce the load on the multi-access system. Later on we may build special hardware, for instance for doing ranging by stereoscopic or focusing methods. In the case of the stereoscopic method we would probably use hardware correlators. We might also build hardware contour followers for the region analysis approach, if it could be shown that a very significant saving in processing time would result.

6. Q) Do you use linguistic methods to recognise the picture input? Is there any trouble when the line drawing of the solids suffers noise? How do you solve the shadow and hidden line problems? What is the most complex solid which your robot can recognise?

A) We are experimenting with a method which involves describing pictures in terms of properties of regions and the relationship between regions (*see* Fennema and Brice, 1969, 1969a). We believe that the system will be moderately immune to noise. The shadow problem will be solved initially by allowing the combination of regions of different intensity level to form a new region and trying recognition again. Later we might attempt to decide whether something was a shadow or not by measuring differences in texture or distance on each side of boundaries between areas of different light intensity.

At present the robot is capable of recognising the simple objects described under heading 1 of this section, under controlled lighting conditions and viewing them from a roughly standard position.

7. Q) Does your robot have colour sensing? What are the merits of this?

A) No. Colour sensing would, however, undoubtedly aid region analysis, and also facilitate communication with the human user concerning a given visual scene. It would be easy to have a single colour-sensitive spot in a moving eye system.

8. Q) How do you solve the difficulties of texture?

A) At present we have no method of coping with texture. In the future we will think of dealing with it by ideas like spatial frequency and spatial correlation, e.g. for distinguishing between textures like wood grain and textures like sand.

9. Q) Which do you think best for range measurements, stereoscopic cameras, range finders as with SRIs robot or sound echo method?

A) Possible methods of range measurement that we are considering are: stereoscopic cameras, focusing adjustment with a monocular camera, and a touch-sensitive probe.

Focusing has the advantage over stereoscopy in that it cannot be deceived by vertical stripes. However it is probably less accurate. We did a little investigation of sound echo ranging techniques but rejected them. The wave-lengths of practical generators are too long for good resolution on our scale of equipment.

10. Q) How does your robot measure a parameter such as size or position of the objects? Are the accuracy and speed of measurement satisfactory for real-time manipulation?

A) At present it does not make such measurements. We are prepared to be satisfied with errors of approximately 5%. Speed limitations are likely to be more severe for vision than for manipulation.

3. Arm and hand

1. Q) Describe the hardware specifications of the manipulators such as degrees of freedom or sensors.

A) A manipulator has been designed and is under construction. Two opposed vertical 'palms' can move independently towards and away from each other over a range of about 18 inches and can move together vertically through about 12 inches. Objects may thus be gripped between the palms, lifted and moved a small distance laterally, in a linear cartesian frame of reference.

Absolute accuracy of positioning will be about 0·2% of full range of movement, but backlash, rigidity and repeatibility should all be only a few thousandths of an inch.

Later, it is intended to add rotation of the manipulator about a vertical axis, and rotation of the palms to turn objects over.

Strain gauges at suitable points will give indications of the forces exerted by the arms and the strength of grip.

2. Q) How dexterous will manipulation be and will it be successful?

A) Too early to say.

3. Q) How do you design the control loop of the manipulators?

   A) The controlling computer will output positional information as 10-bit digital words. These will be converted to an analogue voltage to control a D.C. servo motor. Potentiometers will be used to measure position and tachogenerators to measure velocity.

4. Q) Do you have any suggestions for a system with two hands which would co-operate in a job with human beings?

   A) Not at this stage in terms of implementation. As an application area we have already mentioned teaching aids for children.

5. Q) Do the manipulators have any reflex actions? Is there any need of a small computer for the exclusive use of the manipulators.

   A) A peripheral loop will stop movement if an unexpected force is sensed by the strain gauges.

   Exclusive use of a satellite computer is not necessary. We shall, however, be using such a machine to pre-process visual information and we will make use of it in controlling reflex movements.

4. Locomotion

1. Q) Is there any great need to use legs instead of wheels?

   A) No.

2. Q) How does the robot direct its position in the real world?

   A) Combination of dead-reckoning with landmark-recognition are possible, and have been examined by simulations.

3. Q) Does your robot have balance-detecting and controlling equipment?

   A) No.

4. Q) What are the application fields of robot-like machines with locomotive ability in the near future?

   A) Mowing lawns! If by 'near future' is meant the next two or three years we do not see commercial applications above a rather trivial level.

5. Communication

1. Q) How does your robot communicate with the digital computer?

   A) The robot communicates with the computer as a peripheral of the Multi-POP time-sharing system, running on an ICL 4130 computer. Communication is via transfers of single 8-bit bytes. The output byte is decoded as a command to sample the picture or drive the motors. The input byte contains the state of the bump detectors and brightness of the picture point. When the satellite is installed, communication will be *via* a high-speed link with the ICL 4130. The robot will be interfaced to the satellite, essentially as it is now to the ICL 4130.

6. Brain

1. Q) What performance and abilities does the brain of your robot have? Does it have self-learning ability?

   A) We have engaged in the past in experiments involving developing various abilities in isolation and have not yet finished building an integrated system using these abilities.

   For instance there is the Graph-Traverser program for problem solving (Doran and Michie, 1966; *see* Michie and Ross, 1970 for an adaptive version.) Boxes and memo functions (*see* Michie and Chambers, 1968, Michie, 1968, Marsh, 1970) for rote-learning; programs for deduction and question-answering (Ambler and Burstall, 1969) and the Induction Engine (Popplestone, 1970). Full learning ability requires what is learnt to be expressed in a language more powerful than simply a sequence of weights, as in Perceptrons or Samuel's Checkers learning program.

2. Q) What can the question-answering system in your robot do?

   A) We have implemented a number of approaches to question-answering. We have theorem-proving programs, which, as Cordell Green (1969) has shown, can be modified for question-answering. We also have a program called QUAC based on relational combinators (Ambler and Burstall, 1969).

3. Q) What would be the best interface between robots and human beings?

   A) The best interface from the human's point of view would be spoken and written natural language, together with the ability to point at things with the robot watching through its television camera. In the immediate future, for research purposes, typewriter and visual display using a flexible command language: e.g. 'imperative mode' POP-2.

4. Q) What is the most difficult problem in future artificial intelligence research?

   A) Possibly the internal representation of the robot's world, which will certainly involve automatic methods for inductive reasoning from a very large mass of (mostly irrelevant) data. It seems to us that, to be usable by the robot for serious planning, internal models must involve both *direct* representations in the form of appropriate data structures, as when a map is used to model a terrain, and *indirect* representations in the form of axiom systems and sentences in a formal language such as predicate calculus. Facts are retrieved from the former by look-up and from the latter by reasoning procedures. What is lacking at present is any general theory concerning the relative economics of these two forms of representation, or any principles for automatic transfers of knowledge from one to the other. We are inclined to think that present work on automation of induction will help in the required direction.

   On the deductive side, we would mention the problem of discovering the relationship between solving a problem by logical inference and solving it by an algorithm (i.e. no redundant inferences made), so that opportunities for reducing an inference process to an algorithm may be automatically detected and exploited.

A certain confluence is now apparent between work on robot cognition and the field known as theory of programming. This is because formal equivalences can be set up between proving that a *plan* will be adequate to bring about a given result in the real world and reasoning as to whether a *program* will compute a given function (*see* Green, 1969a). We attach importance in this connection to recent advances in the theory of formal proofs about programs (Floyd, 1967; Manna and McCarthy, 1970; Burstall, 1970).

In terms of implementing systems capable of operating within reasonable time constraints, methods for handling highly parallel processes will be crucial, and these are still in their infancy.

Acknowledgements

The Edinburgh work has on the hardware side mainly been conducted in the Bionics Research Laboratory (of the Department of Machine Intelligence and Perception) founded by Professor R. L. Gregory. The basic software and interfacing work has been done in the Experimental Programming Unit of the same Department. The authors have been greatly assisted by their colleagues on the Robot Working Party: Dr. R. M. Burstall, Reader in Machine Intelligence, Dr. J. A. M. Howe, Director of the Bionics Research Laboratory and Dr. H. R. A. Townsend, Senior Consultant Neurologist in the Western General Hospital and part-time Senior Lecturer in the above Department.

A special debt is acknowledged to an overall philosophy contributed by Professor R. L. Gregory. His dictum: 'The cheapest store of information about the real world is the real world itself' was a major part of the original motivation, and the emphasis laid upon the role of internal models in Gregory's analysis of perception (*see* Gregory, 1966) continues to be central to our work.

References

AMBLER, A. P., and BURSTALL, R. M. (1969). Question-answering and syntax analysis, *Experimental Programming Reports, No.* 18, Department of Machine Intelligence and Perception, University of Edinburgh (March, 1969).

BURSTALL, R. M. (1970). Formal description of program structure and semantics in first order logic, *Machine Intelligence 5*, Meltzer B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 79-98.

BURSTALL, R. M., and COLLINS, J. S. (1971). A primer of POP-2 programming, *Programming in POP-2*, Burstall, R. M. (Ed.), Edinburgh University Press, Edinburgh (in preparation).

BURSTALL, R. M., and POPPLESTONE, R. J. (1971). POP-2 Reference Manual, *Programming in POP-2*, Burstall, R. M. (Ed.), Edinburgh University Press, Edinburgh (in preparation). Also in *POP-2 Papers*. Edinburgh University Press, Edinburgh (1968).

DORAN, J. E., and MICHIE, D. (1966). Experiments with the Graph Traverser Program, *Proceedings of the Royal Society, A.*, Vol. 294, pp. 235-259.

FEIGENBAUM, E. A. (1968). Artificial intelligence: themes in the second decade, *Proc. IFIP Congress 1968*, North-Holland Publishing Co., Amsterdam, pp. 10-24.

FENNEMA, C. L., and BRICE, C. R. (1969). A region oriented data structure, *Technical Note No. 7, S.R.I. project 7494*, Stanford Research Institute, Stanford.

FENNEMA, C. L., and BRICE, C. R. (1969a). Scene analysis of pictues using regions, *Technical Note No. 17, S.R.I. project 7494*, Stanford Research Institute, Stanford.

FLOYD, R. W. (1967). Assigning meanings to programs, *Mathematical aspects of computer science*, Amer. Math. Soc., Providence, Rhode Island, pp. 19-32.

GREEN, C. C. (1969). Theorem proving by resolution as a basis for question-answering systems, *Machine Intelligence 4*, Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 183-205.

GREEN, C. C. (1969a). Applications of theorem-proving to problem-solving, *Proceedings of Int. Joint Conf. on Artificial Intelligence*, Walker, D., and Morton, L. M. (Eds.), Washington D.C., pp. 219-239.

GREGORY, R. L. (1966). *Eye and Brain*, Weidenfeld and Nicolson.

MANNA, Z., and MCCARTHY, J. (1970). Properties of programs and partial function logic, *Machine Intelligence 5*, Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 27-37.

MARSH, D. L. (1970). Memo functions, the Graph Traverser and a simple control situation, *Machine Intelligence 5*, Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 281-300.

MICHIE, D. (1968). "Memo" functions and machine learning, *Nature*, Vol. 218, pp. 19-22.

MICHIE, D., and CHAMBERS, R. A. (1968). BOXES: an experiment in adaptive control, *Machine Intelligence 2*, Dale, E., and Michie, D. (Eds.), Oliver and Boyd, Edinburgh, pp. 137-152.

MICHIE, D., and ROSS, R. (1970). Experiments with the adaptive Graph Traverser, *Machine Intelligence 5*, Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 301-318.

POPPLESTONE, R. J. (1970). An experiment in automatic induction, *Machine Intelligence 5*, Meltzer, B., and Michie, D. (Eds.), Edinburgh University Press, Edinburgh, pp. 203-215.