CONTROL ENGINEERING AND ARTIFICIAL INTELLIGENCE

B.R.Gaines

Department of Electrical Engineering,
University of Essex, Colchester, Essex, U.K.

1 Background

The relationship between control engineering (CE) and artificial intelligence (AI) is significant for workers in both disiplines:-

SYSTEM-THEORETICALLY there is much common ground in both objectives and activities, for example: 'identification and modelling' in CE cf 'knowledge acquisition and structuring' in AI, or 'control' in CE cf 'goal-seeking' in AI. By way of contrast, however, 'stability and sensitivity' analysis so central to CE seem to have no real equivalent in AI (perhaps some concept of 'robustness' of knowledge structures and of strategies), and the linguistic information processing so ubiquitous in AI has no real equivalent in CE (much to its detriment - see section 4).

TERMINOLOGICALLY both disiplines have borrowed their vocabularies freely from human analogies, e.g. 'adaptive and learning' controllers and robots, and this leads to a superficial resemblance that does not necessarily denote a true correspondence and may be misleading.

SOCIOLOGICALLY both disiplines are new and have had to establish their identities in the face of substantial opposition. Perhaps CE has had the easier time of it in growing out of well-established Engineering Departments, whereas AI has had to struggle free of Computer Science, itself struggling free from Mathematics Departments. However that is another story not appropriate to these notes - it is worth noting that the problems of AI as a subject area are not unique but occur with most new developments (see Merton 1973 for a deep and fascinating account of such aspects of the sociology of science).

I cannot pretend in these notes to give a thorough account of CE and its relationship to AI. They are intended as a guide to the mainstream of activities in CE and to some real, and some apparent, points of contact with AI. The literature references are deliberately highly selective to make it feasible and worthwhile to follow them up. The next section outlines the main subject areas of control theory; the following section that of control application; and the final section attempts to sum up the current state of the art.

2 Control Theory

The theoretical foundations of CE may be broadly divided into a logical sequence of three main subject areas:

(a) System Identification - from the input/output behaviour of a system determine a suitable model for its structure - in CE the 'grey-box' approach is most common in that we assume quite a lot about the structure of the system to be modelled (generally that it is linear and not very complicated) in order to acquire a model fairly rapidly - if the model proves a poor predictor this shows up at the development stage and the assumptions are changed (perhaps this is the story of science !) - the main complication is that we can rarely indulge in identification alone, but either want to, or have to, control also to some extent (because it would blow up, or otherwise not be the same system if we did not), - this leads to the complications of the so-called 'dual control problem' of learning about a system whilst we are at the same time using that knowledge to control it - this problem is truly fundamental and may be studied in its simplest form as the 'two-armed bandit problem' (Witten (1975) is an excellent summary of these problems - Eykhoff's (1974) book summarizes well the main contributions to system identification);

(b) Control itself is clearly central - the significant factors are that it implies a purpose and hence also a means of determining when, or how well, that purpose has been achieved - the classical picture of control is that we use this measure of success to determine the control action - 'negative feedback' is a powerful paradigm that the cybernetician's are quick to see everywhere - however it is significant that much, perhaps most, industrial control does not employ such short-term negative feedback of its primary performance variables - 'open-loop' control is very common in which a plant is taken through a standard cycle without any on-line evaluation of the results - negative feedback is used to control short-term secondary variables like the temperatures of various sub-systems, but performance feedback is by off-line chemical analysis of the products and its effective is retrospective,strategic rather than tactical - in this respect an industrial plant strongly resembles a living organism with its local feedbacks for posture control underpinning its capability for action (in the philosophical sense - Care & Landesman 1968) but the feedback on the result of actions being a long-term process of cogitation and learning. The moral is, "don't take too much notice of the simple diagrams of negative feedback at work in the control textbooks" - they are the spinal cord of control, not its brain !.

(c) Stability and Sensitivity analyses come closest to being what control theory is really about - having got a model for a system and based the design for a controller upon it then comes the grand questions, 'how will the overall system respond to perturbations of the controlled variables' (stability analysis) and 'how will the system performance change if the model is not quite right' (sensitivity analysis) - it was the stability analysis techniques of Bode and Nyquist in the 30's that made control engineering a subject area - later progress may be measured in terms of increasing generality of such analyses (e.g. Lyapunov's criteria - Leondes 1965) - and the weak areas of CE today are those where such criteria are inapplicable, i.e. most nonlinear systems. To my mind the development of techniques for the study of stability and sensitivity (general 'robustness') is the major determinant of the success of any field of applied system theory, including AI. We need to know the potential effects of defects in our own knowledge - in practice the best technique is not necessarily that which works best under prescribed conditions, but that which has limited 'downside potential' under a wide range of possible conditions (Radanovic (1966) gives a useful overview of the role of stability and sensitivity analyses in control theory).

2.1 Linear Systems Theory

If I were fool enough to try to define the differences between the theoretical sub-structures of CE and AI in one sentence, it would be that, "CE theory is based upon modelling continuous processes as linear systems, whereas AI theory is based upon modelling discrete processes as linguistic systems". The role of linear systems theory in CE cannot be over-emphasized - it is both the major strength of the whole subject area, and its major weakness. Linearity is a powerful constraint that leads to deep and elegant results - that they are applicable to physical systems is metaphysical magic arising from: (a) Mathematically, linearity is a good approximation to the local behaviour of a continuous system with bounded higher order derivatives (since these correspond to energy transfers they tend to be bounded in all but bombs); (b) We chose the systems to be controlled, usually manufacturing them, so that they may be selected to be linear; (c) We certainly chose the controllers and, with the application of sufficient negative feedback, a linear controller tends to make the whole system appear linear.

That, perhaps, is a sufficient rationale for the domination of CE by linear systems theory and, re-interpreted, explains the lack of impact of such theory in AI where we rarely deal with the dynamics of physical systems and choose our problem areas with different theoretical models in mind. This great divide, although very real, is a historical one that cannot last for ever - control engineers have been turning their attention to biological and economic systems where the assumption of linearity is only partially successful, and AI robots are increasingly facing the nasty, continuous physical world. Hendrix's (1973) AI paper on "modelling simultaneous action and continuous processes" states the problem clearly and bridges the gap between the two disiplines. It may be contrasted with CE papers on computer modelling of physical systems (Lee et al 1968).

It came as a shock to control engineers in the 50's to discover that the optimal controller for a linear system was not itself linear. Pontryagin's famous 'maximum' principle (Leondes 1965) showed that to change the state of a linear system in minimum time one always had to apply the maximum available force to it. Only the sense of that force varied and this was determined by a 'switching surface' in the state space of the system. This resulted in a massive effort to determine the forms of such 'optimal controllers' for a wide variety of systems. It was a major change of direction for CE because the switching surfaces made the control algorithm a discontinuous, decision-making strategy, utterly different from that of classical continuous control. It was soon discovered (Wilde & Wescott 1962) that the human controller used such 'bang-bang' strategies in limb movements so that, as so often happens, nature had been there all the while. The application of optimal control to industrial systems became feasible with the advent of on-line computer control but, and here comes the moral of the story, it proved very disappointing in practice largely because the optimal controllers tended to be extremely sensitive to errors in plant models, becoming very non-optimal when the real world was not precisely as expected. Fuller's (1967) classic paper showed that the old linear controller was not all that sub-optimal, and far more robust to changing plant characteristics. In retrospect one can see that the definition of optimality was at fault, and that in real CE the goal is not the optimal control of one prescribed plant, but the reasonable control of any of a range of possible plants. This concept was given a rigorous interpretation by Kwakernaak (1965) who termed a controller 'admissible' for a range of plants if their was no other controller that was uniformly better for all of them (a concept borrowed from statistics, and a very useful one).

2.2 The Notion of 'State' and Automata

There is one important concept largely developed in CE which we probably all take for granted so much as to forget its comparatively recent origins. The 'state' of a system may be defined as the minimum information necessary to account for its potential future behaviour. Whilst the concept is a straighforward one in terms of system structure, it becomes very much more subtle when used to define the 'state' of a model of a system determined from its behaviour. Many systems theorists in recent years have analysed the concept of "state" in ways that throw interesting light on tha nature of explanation and causality - Zadeh (1964), Zadeh & Desoer '1963) or Zadeh & Polak (1971) are good introductions.

In a recent article, Arbib (1975b) points out that Wiener, who as the father of cybernetics might be expected to be very state-orientated, did not use the concept - his was a more "rheological" approach relating the entire past history of a system to its future behaviour. Whilst modern control theory may be seen to be crucially based on the notion of system state, this is not necessarily the most appropriate basis for system development or modelling, particularly of biological systems. For example, the "causal" explanation of the behaviour of an animal as determined by its current internal state and external input may be extremely complex, whereas the "teleological" explanation in terms of its intended future behaviour may be very simple - Rescher (1963) gives some interesting examples of this in terms of automata. States have a definitional importance that makes them appear fundamental, but it is important to realize that the "inferred state" is an artefact of the observing system and does not necessarily have an analogy in the structure of the observed system.

Automata theory is the ultimate expression of general state-determined behaviour and it has its place in most modern textbooks on control theory. Its actual role so far has been peripheral to both control theory and practice, largely because of the lack of powerful general results comparable to those of linear systems theory. Arbib (1975a) has recently provided a useful critique of the role of automata theory in modelling biological systems, and, over a span of some years, he has been active in promoting the role of automata theory in control engineering (Arbib 1965, 1966, Arbib & Zeiger 1969) and its relationship to linear systems theory (Arbib & Manes 1974). This series of papers is well worth reading both for its general interest and as an introduction to algebraic techniques in CE, particularly the work of Kalman (et al 1969).

3 Control Applications

Whilst the devlopment of control theory has its lessons for, and interactions with, AI, it is surely in the applications, and attempted applications, of CE that most value is to found. Finding suitable problems for the application of AI techniques is a continuing requirement (where 'suitable' implies both relevance and borderline solvability), and advanced CE may be seen as defining the lower end of a spectrum of which realistic AI defines the other (standard CE is the infra-red, and optimistic AI the ultra-violet, of this spectrum !).

By far the best sources of information on control applications are the symposia of the International Federation for Automatic Control (IFAC) and its associated journal, 'Automatica'. Many control journals are more concerned with applied (?) mathematics than with possible applications, but IFAC has the well-deserved reputation of keeping theory closely allied to reality and of treating experience in the complex problems of controlling real systems with respect, even if it does not necessarily illustrate elegant general systems principles.

The topics of the 12 IFAC symposia planned for 1976 include: control in space; automation in paper, rubber and plastics; automation in offshore oil field operation; dynamic modelling and control in national economies; control mechanisms in bio- & eco-systems; etc. The range of interests is very wide and barely hinted at by most textbooks on CE. The proceedings of these symposia, and those of the IFAC congresses held once every 3 years (latest August 1975, Boston, USA), should be the prime sources for those interested in finding interesting CE situations in which to apply AI techniques.

4 Concluding Remarks

In the 60's there appeared to be much common ground between AI and CE - "learning controllers" in particular seemed to represent AI thinkinf in CE (see Tsypkin 1971 for a review of this type of work, and Gaines 1969 for a discussion of the vocabulary of 'adaption' in biology and engineering). In the 70's, however, this communality has lessened with the decrease of interest in machine learning as such in AI, and the disenchantment with the more esoteric forms of "learning machine" in CE (and the absorption of those that worked into the less emotive vocabulary of 'optimization', etc.). Looking through the first 5 years of 'Artificial Intelligence' I find nothing which can be taken as a CE application. Looking through the programme for the IFAC 75 congress there is little that can be directly related to AI.

However, this apparent gulf is not as real as it appears since many workers in CE are fully aware of the AI literature, and vice versa. Fundamentally the problems treated in CE, particularly now that economic and biological systems are attracting such attention, constitute a natural domain for AI - they are problems where human activities are being aided or replaced by automated systems, and the frontiers of this work must, virtually be definition, be closely allied with AI.

Finally, to end on a positive note, let me illustrate the interesting developments that may arise when AI techniques are applied to a real control system by describing some recent work at Queen Mary College (Assilian 1974, Mamdani 1974, Assilian & Mamdani 1975). A model steam engine has been used as a test bed for various forms forms of 'learning controller' - chosen because of its nonlinearities, hysteresis, holdups, long time-constants, and other real-life features (like proneness to blow-up !). The initial objective was to investigate a range of published learning algorithms comparing them amongst themselves and with some classical linear controllers. It was realized that the learning controllers on any real plant would, in their initial naïve state, close it down, or destroy it, before acquiring much information about it, and attention was focussed upon the possibility of 'priming' the controller with a reasonable initial strategy.

Zadeh's (1973) application of fuzzy sets theory to linguistic reasoning seemed to offer an opportunity to do the 'priming' in a general form within the spirit of an AI approach by giving the controller initial "instructions" of the form, 'if the pressure error is negative big, and the change in pressure error is near zero then increase the heat input slightly'. A fuzzy control strategy of this form was constructed and used quite formally by the controller as a computer program embodying the rules of fuzzy logic. The resultant performance was better than that of the best linear controllers and way beyond that ever achieved by any of the 'learning machines' !

The most fascinating aspect of this story is the way in which the synthesis of a control policy using fuzzy reasoning so mirrors one's intuitive grasp of how controllers are developed. The design engineer 'plays' with the plant, getting the 'feel' of it before taking formal measurements. His model of the plant and the instrumentation to enable this to be quantified are conditioned by his knowledge of it, and those like it, and by his own interaction with it. Currently this knowledge is translated into the mathematics of linear system theory and the design of a controller then becomes a mathematical problem. The direct linguistic transformation of statements about the plant into a control strategy is in many ways more natural. Mamdani has recently produced independent evidence that this is so: Chapter 10 of Perry & Waddell's (1972) book on rotary cement kilns is an operations handbook for the human operator of a lime kiln. Since it is intended to be vàlid for any kiln it is couched in qualitative rather than numeric terms, and consists of 27 rules of the form: "if the burning zone temperature is drastically low, the percentage of oxygen is low and the back-end temperature is low then reduce the kiln speed and reduce the fuel", etc., a 'fuzzy' control strategy priming the human controller !

Like algebraic techniques, or any innovation in any subject area, it will take time for results such as those of Mamdani & Assilian to be tested in a range of situations and assimilated into CE. However, they do illustrate that the gap between the continuous linear world of CE and the discontinuous linguistic world of AI is not insuperable.

5 References

M.A.Arbib (1965) A common framework for automata theory and control theory, SIAM Journal of Control, vol.3, pp.206-222.

M.A.Arbib (1966) Automata theory and control theory: a rapprochement, Automatica, vol.3, pp.161-189.

M.A.Arbib (1975a) From automata theory to brain theory, International Journal of Man-Machine Studies, vol.7, pt.3.

M.A.Arbib (1975b) Cybernetics after 25 years: a personal view of system theory and brain theory, IEEE Transactions on Systems, Man and Cybernetics, May 1975, vol.SMC-5, pp.359-363.

M.A.Arbib, and E.G.Manes (1974) Foundations of systems theory: decomposable machines, Automatica, vol.10, pp.285-302.

M.A.Arbib, and H.P.Zeiger (1969) On the relevance of abstract algebra to control theory, Automatica, vol.5, pp.589-606.

S.Assilian (1974) Artificial intelligence in the control of real dynamic systems, PhD thesis, Queen Mary College, University of London, UK.

N.S.Care, and C.Landesman (1968) Readings in the Theory of Action, Bloomington: Indiana University Press.

P.Eykhoff (1974) System identification, London: Wiley.

A.T.Fuller (1967) Linear control of nonlinear systems, International Journal of Control, vol.5, pp.197-243. Also reprinted in a useful collection: A.T.Fuller (ed) (1970) Nonlinear Stochastic Control Systems, London: Taylor & Francis.

B.R.Gaines (1972) Axioms for adaptive behaviour, International Journal of Man-Machine Studies, vol.4, pp.169-199.

G.G.Hendrix (1973) Modelling simultaneous actions and continuous processes, Artificial Intelligence, vol.4, pp.145-180.

R.E.Kalman, P.L.Falb, and M.A.Arbib (1969) Topics in Mathematical Systems Theory, New York: McGraw Hill.

H.Kwakernaak (1965) Admissible adaptive control, Proc. IFAC Symposium on the Theory of Self-Adaptive Control Systems, Teddington, UK.

T.H.Lee, G.E.Adams, and W.M.Gaines (1968) Computer Process Control: Modeling and Optimization, New York: Wiley.

C.T.Leondes (ed) (1965) Modern Control Systems Theory, New York: McGraw Hill.

E.H.Mamdani (1974) Application of fuzzy algorithms for control of simple dynamic plant, Proc. IEE, vol.121, pp.1585-1588.

E.H.Mamdani, and S.Assilian (1975) An experiment in linguistic synthesis with a fuzzy logic controller, International Journal of Man-Machine Studies, vol.7, pp.1-13.

R.K.Merton (1973) The Sociology of Science, Chicago; University of Chicago Press.

K.E.Perry, and J.J.Waddell, The Rotary Cement Kiln, New York: The Chemical Publishing Co.

L.Radanovic (1966) Sensitivity Methods in Modern Control Theory, Oxford: Pergamon Press.

N.Rescher (1963) Discrete state systems, markov chains, and problems in the theory of scientific explanation and prediction, Philosophy of Science, vol.30, pp.325-345.

Y.Z.Tsypkin (1971) Adaption and Learning in Automatic Systems, New York: Academic Press.

R.W.Wilde, and J.H.Wescott (1962) The characteristics of the human operator engaged in a tracking task, Automatica, vol.1, pp.5-19.

I.Witten (1975) The apparent conflict between estimation and control - a survey of the two-armed bandit problem, Technical Report EES-MMS-BAN-75, Department of Electrical Engineering Science, Essex University, Colchester, UK.

L.A.Zadeh (1964) The concept of state in system theory, in M.Mesarovic (ed) Views on General Systems Theory, New York: Wiley, pp.39-50.

L.A.Zadeh (1973) Outline of a new approach to the analysis of complex systems and decision processes, IEEE Transactions on Systems, Man, and Cybernetics, vol.SMC-3, pp.28-44.

L.A.Zadeh, and C.A.Desoer (1963) Linear System Theory, New York: McGraw-Hill,

L.A.Zadeh, and E.Polak (eds) (1969) System Theory, New York: McGraw Hill.