

An Engineering Model of Subcognition*

Donald Michie

TR 92-22

Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, Virginia 24061

May 14, 1992

*This report is cross-listed as Statistics Department report 92-8.

Dr. Donald Michie is the Director for Advanced Study at The Turing Institute, 36 North Hanover Street, Glasgow G1 2AD, UK. This department is honored to host Dr. Michie as the C. C. Garvin Endowed Visiting Professor in Computer Science and Statistics for spring semester 1992.

AN ENGINEERING MODEL OF SUBCOGNITION

Donald Michie
The Turing Institute,
Glasgow, UK
May, 1992

ABSTRACT

Symbolic and neural models are discussed from the point of view of implementing the "intuitive" processes which mediate real-time control skills. The engineering preference is to transcend rather than emulate those limitations which are inherent in the brain's subsymbolic modes of operation. A main payoff lies in recovering, from performance data, symbolic representations of intuitive real-time skills. Practical illustrations are reviewed using recent case studies.

INTRODUCTION

The labels "strong AI" and "weak AI" have sometimes been used to differentiate two schools. Criteria are summarised below.

	Strong	Weak
Feasibility of goals	Human-level intelligence will be achieved in machines within foreseeable time.	Human-level intelligence will be implemented only in some unimaginable future, or perhaps never.
Forms of implementation	All thought can be mechanized as sequential logical reasoning from axiomatic descriptions of the world. The "physical symbol system hypothesis": all agents, including intelligent, are best implemented symbolically.	Most thought is intuitive. not introspectable, non-logical, associative, approximate and "fuzzy". best modelled by brain-like ultra-parallel networks.
Personnel	Vintage AI professionals, e.g. Turing, Simon, Newell, Minsky, Feigenbaum, Nilsson, and their followers.	Members of other professions, particularly in linguistics, neurobiology and philosophy.

The above taxonomy lays emphasis on the "physical symbol system hypothesis" of Newell and Simon (1976). Their intended interpretation restricts symbol systems to those which can transparently support communication with human users. Thus the lists of numerical weights in which neural nets express themselves constitute "symbols" of a sort, but not in the sense intended by the above authors. This restriction has persuaded some practitioners that the physical symbol system hypothesis entails exclusion of intuitive processes from AI's domain of discourse. Such separatism is unsafe, since much knowledge-based thought seems irredeemably intuitive and sub-articulate (for a recent commentary see French, 1990). For its subcognitive processes, there is no direct evidence that the brain employs a symbolic regime. Hence those who accept subarticulate expertise as a proper AI concern may wonder whether for this purpose they should abandon symbolic representations as untrue to nature. The present review advocates a different position, namely that a conceptually transparent symbolic style offers a way of improving on nature. By representing intuitive processes symbolically, inductive inference can do something which is both non-brainlike and also highly useful, catering to the client who says: "My in-house experts may be 'intuitive'. But I want an expert system to formulate its reasons more explicitly than that."

KNOWLEDGE AND THOUGHT

In industrial knowledge systems the implementer has to distinguish between thought as something to be communicated and thought as problem solving. Choice of representation remains a developer's option. In implementing intuition, he or she may decide that it is something over which to draw a veil. The veil may be woven of neural nets, or of hand-crafted spaghetti-code, or of something else. But suppose that the developer has to supply the customer also with means to draw the veil aside, — for purposes of interrogation about goals, plans, evidence, justification and the like. At the price of

being less true to nature, he or she might then be better off not to have veiled it in the first place. Like cognitive and brain scientists, knowledge engineers also study the structure of expertise. *Unlike* cognitive and brain scientists, they do this (or should do) for the purpose not of emulating but of transcending the brain's limitations. First among these is the relative inarticulacy of what both cognitive scientists and knowledge engineers call "procedural knowledge", thus distinguishing it from "declarative".

Declarative Knowledge

It is characteristic of the retrieval and use of declarative knowledge that it is ordinarily done in conscious awareness. From a wealth of neurobiological observations concerning the effects of brain lesions on memory, L. R. Squire (1987, chapter 11) distinguishes declarative memory from procedural as "memory that is directly accessible to conscious recollection". By contrast, the hall-mark of a highly trained expert brain is that it does much of its work intuitively. "Dialogue elicitation" of rules for building expert systems may therefore be frustrated whenever a given expertise involves strategies stored in procedural memory. Inaccessibility to consciousness of even parts of a targeted expertise can then cause serious problems for large knowledge engineering projects, such as Japan's ambitious "Fifth Generation" (Michie, 1988). Differentiation of the two forms is thus desirable.

Declarative knowledge comprises whatever lends itself to logical formulation: goals, descriptions, constraints, possibilities, hypotheses. The declarative category also includes *facts*. When these relate directly or indirectly to events in the agent's own experience, their place of storage is referred to as "episodic" memory. Another subdivision of declarative knowledge is held to reside in "semantic" memory, which Squire defines as follows:

Semantic memory refers to knowledge of the world. This system represents organized information such as facts, concepts, and vocabulary. The content of semantic memory is explicitly known and available for recall. Unlike episodic memory, however, semantic memory has no necessary temporal landmarks. It does not refer to particular events in a person's past. A simple illustration of this difference is that one may recall the difference between episodic and semantic memory, or one may recall the encounter when the difference was first explained.

A school founded by John McCarthy (1959) aims to extend formal logic to serve as a vehicle for mechanizing declarative knowledge (*see* a recent collection edited by Ginsberg, 1987). I will say little further about the project, beyond expressing my respect for such work. Its philosophical importance is matched only by its difficulty. My theme is closer to the name and nature of expert systems. These are not so much to do with endowing computers with knowledge of the world, as with equipping them with useful know-how. In face of the difficulties which confront the McCarthy project, there is something to be said for studying the mechanization of procedural knowledge as a separate problem.

Nature of Procedural Knowledge

In Anderson's (1990) text on cognition, skilled procedures are pictured as arising in part by derivation from pre-existing mental descriptions. No direct evidence is offered. Knowledge engineers concerned with real-time skills have been led by practical experience in a rather different direction. The empirical picture is one of inductive compilation from sensorimotor data gathered in the course of trial and error. In this picture the role of higher-level knowledge is not to participate directly, but to steer the learning process, setting and adjusting the frame within which skill-bearing rules are constructed.

The final phase of skill-learning, described by Anderson and others as “automatization”, does not ordinarily support introspective report by the expert performer, — hence the “knowledge-acquisition bottleneck” of applied AI. Procedural knowledge, as we have seen, limits itself to the “how to” of skilled tasks, whether physical as in making a chair, or more abstract as in prediction of sterling rates against the dollar or the differential diagnosis of acute abdominal pain. A common synonym for such knowledge is know-how, and its manifestation in observable behaviour is called “skill”. One difficulty is that observed task-performance does not necessarily reveal whether a given expert’s behaviour really exemplifies a skill in the procedural sense or whether he or she is using declarative-semantic memory to form action-plans on the fly. Squire’s earlier-cited definition supplies a test, namely the ability to render a verbal account of the way in which each decision was made, — possible only for declarative memory. A second criterion is the frequency of the recognise-act cycle: this may simply be too fast for mediation by “what-if” inferential planning to be feasible.

For those concerned to recover procedural rules, as in building expert systems, the lack of verbal accessibility (on which Anderson also remarks) should give pause. Yet there is widespread faith among knowledge engineers that special methods of “dialogue elicitation” can be found which will permit the construction by this method of rule-based systems on the scale of such inductively built systems as the GASOIL (Slocombe et al., 1986) and BMT programs of Table 1.

	APPLICATION	NO. OF RULES	DEVELOP. MAN-YRS	MAINTENANCE MAN-YRS/YR	INDUCTIVE TOOLS
MYCIN	medical diagnosis	400	100	N/A	N/A
XCON	VAX computer configuration	8,000	180	30	N/A
GASOIL	hydrocarbon separation system configuration	2,800	1	0.1	ExpertEase and Extran 7
BMT	configuration of fire-protection equipment in buildings	>30,000	9	2.0	1st Class and RuleMaster

Table 1 Of the world’s three largest expert systems the two latest (GASOIL and BMT) were not constructed from rules obtained in dialogue fashion, but by automated induction from expert-supplied data. In each case the induction engineer trained the system in the desired skill in the style that the master of a craft trains an apprentice, — by a structured sequence of selected examples. Rates of code production are typically in excess of 100 lines of installed Fortran, C, Pascal etc. per programmer day. The methodology allows validation to be placed on a user-transparent basis (Michie 1989), and maintenance costs are in many cases trivialised. Tabulation is from Slocombe et al. (1986) with 1990 data on BMT added. The BMT program is described on p.10 of *Pragmatica* vol. 1 (ed. J.E. Hayes Michie), Glasgow, UK: Turing Institute Press.

Is rule induction from expert-supplied data nevertheless in some sense a second-best option for building systems on the BMT scale? The answer hangs on the observed results of allowing the experts to communicate their skills (as in the BMT case) solely *via* illustrative responses to selected cases. Does he or she thereby omit something indispensable? Certainly the practitioner's explicit and communicable awareness is integral to expertise in *some* task domains. But other domains, which lack this property, can be found not only among a rather wide variety of industrial tasks, but even in such purely "mental" forms of expertise as playing a strong game of checkers (see below).

As a paradigm of procedural knowledge, Feigenbaum and McCorduck (1983, p.55) give the example of tying one's shoes. It is interesting that once this skill has reached the stage known as automatization it can continue unaffected by massive disruption of the individual's brain mechanisms for acquiring and handling important forms of declarative knowledge. Damasio describes a patient named Boswell. The following summary is from Patricia Smith Churchland (personal communication).

In addition to losing the hippocampal structures, he has massive damage to frontal cortex. He can identify a house, or a car, but he cannot identify *his* house or *his* car; he cannot remember that he was married, that he has children, and so forth. He seems to have no *retrograde* episodic memory, as well as no *anterograde* episodic. ... Boswell can still play a fine game of checkers, though when asked he says it is bingo. He cannot learn new faces and does not remember "pre-morbid" faces such as that of his wife and his children... Boswell can play checkers, tie his shoes, carry on a conversation, etc.

Of considerable interest is the survival of Boswell's checkers skills. Evidently what we shall later term "fast" skills are not the only ones for which procedural knowledge may be so dominant as effectively to eclipse declarative. In contrast to chess skill, checkers was already known not to lend itself to the planning approach and to be essentially "intuitive". When A.L. Samuel was engaged in his classic studies of machine learning using the game of checkers, he had numerous sessions with leading checkers masters directed towards dialogue acquisition of their rules and principles. Samuel reported (personal communication) that he had never had such frustrating experiences in his life. In terms of discernible relationship to what the masters actually did, the verbal material which he elicited contained almost nothing which he could use or interpret. In similar vein, Feigenbaum and McCorduck (*loc. cit.*, p.82) describe this type of expert response in the following terms: "That's true, but if you see enough patient/rocks/chip-designs/instrument readings, you see that it isn't true after all." They conclude "At this point, knowledge threatens to become ten thousand special cases."

The message from clinical studies is that skilled performance of even sophisticated tasks can still be manifested, and learned, when the brain is so damaged that knowledge of new happenings cannot be retained and previously stored facts and relations (declarative-semantic memory) are seriously disrupted. Another circumstance under which the mediation of declarative memory is at least equally disabled can be observed in the normal brain by imposing a sufficiently restrictive constraint on the time available for the recognise-act cycle, as in touch-typing. This skill does not depend on the storage and retrieval of declarative knowledge, and can be acquired and executed in its virtually complete absence. Recall that when copy-typing at speed the typist does not need to understand the words as he or she reads them. Indeed, after a speed test little or nothing of the text's content can be recalled. Moreover, educated onlookers are surprised, although they should not be, by the outcome of a request to the typist (supposing that he or she has been using a typewriter with unlabelled keys) to label the keyboard correctly with the proper alphanumeric symbols. Lacking a declarative model, the touch-

typist is ordinarily unable to do so (see, for example, Posner, 1973), other than by the expedient of deliberately typing a symbol and observing where the finger went!

Simon (in press) has recently re-emphasised that simple recognition of a familiar object takes at least 500 milliseconds. Operations involving reference to a semantic model of the task domain require retrieval from long-term memory of relatively complex knowledge-structures and an associated apparatus for inferring, storing and utilizing intermediate results. Such elaborate transactions are to be found only in the "slow lane". Here seconds, minutes or even hours are required to incubate a decision. The bare bones of an explicit rationale for a slow-lane decision, when it comes, can usually be elicited from the expert by verbal report. Not so in the fast lane, to which the present discussion is confined. "Fast" skills cannot be accessed by "dialogue elicitation" methods. How then are expert systems to be built for these skills? A way round the seeming impasse is to record behavioural traces from the expert subject. Inductive inference then reconstructs from recorded decision-data rule-based models of the brain's hidden strategies. As reported in this review, machine execution of data-derived models has been found to generate performance exceeding in reliability the trained subjects' own.

Postulates of Skill Acquisition

Experimental work which will now be described was animated by a point of view about brains, summarised below as a list of postulates. Declarative knowledge is abbreviated to "D" and procedural to "P". To avoid committing to a view about embryonic forms, P designates only procedural knowledge which has already reached the automatized stage.

1. human agents are able verbally to report their own D;
2. human agents cannot verbally report their P;
3. D can be augmented by being told, and also by deduction;
4. P is built by learning, whether by imitation or by trial and error;
5. P can be executed independent of D, but not *vice versa*;
6. decision-taking *via* P is fast relative to use of D;
7. sufficiently fast control skills depend on P alone;
8. even for some slow skills P is sufficient for expert performance;
9. rule-induction can extract an explicit form of P from behavioural traces.

Experiments on dynamical control have yielded illustrations of the listed postulates, culminating in a test of 9 above, namely induction of rules from silent brains. But a comment is first requisite on the undoubted existence of expert systems (EXCON was mentioned earlier) whose rule-bases have, with whatever difficulty, been constructed by dialogue acquisition.

Many observers have noted that experts seek to escape from the requirement of rule-formulation (which they find uncongenial) by supplying "rules" of such low-level form that they constitute no more than concocted sample cases, i.e. specimen decision-data. The phenomenon has been described by Sterling and Shapiro (1986) in their description of the construction of a credit evaluation expert system. The finance specialists continually gravitated towards concrete instances rather than general rules. This has indeed been a universal finding in knowledge engineering, in line with the known facts concerning procedural memory and its mode of access.

But what if knowledge engineers in search of improvements on raw formulations, were consciously or unconsciously to apply their own powers of inductive inference to such sample cases? They could then themselves create the kind of high-level rule structures that they had hoped to elicit. The result would of course be testimony more to their own powers of inductive generalization than evidence that

experts can introspect their own rules. In a recent aerospace application two knowledge engineers were able, by deliberately exploiting this style of "rule-conjecture and test", to construct a rule-based solution with no more than a black-box simulator of the task domain to provide corrective feed-back. No set of rules pre-existed, either in an expert's brain or anywhere else.

AN EXPERIMENT IN RULE-BASED CONTROL

The role of the systems developer postulated above requires only a reactive oracle from which he or she can elicit verdicts. This source need not be an expert. Indeed, it need not be human. As will be described it could be a simulator on which the developers can play "what-if" games with their latest conjectured rules ("what if we modify the rules like this? ... what would result from that adjustment? ... etc."). In an R & D contract for a US space consortium (Sammut and Michie, 1990), Dr. Claude Sammut and I were given access to just such an interactive oracle.

When building a controller for a physical process, traditional control theory requires a mathematical model to predict the behaviour of the process. Many processes are either too complicated to model accurately or insufficient information is available about the process environment. Space-craft attitude control is an example of the latter. Our client was interested in the development by machine learning of a rule-structured controller. A check was desirable as to whether dynamical control tasks can be satisfactorily handled by production rules at all, whether these are captured by learning algorithms or developed in some other way.

If the attitude of a satellite in low Earth orbit is to be kept stable by means of thrusters, the control system must contend with many unknowns. For example, although very thin, the Earth's atmosphere can extend many hundreds of kilometres into space. At different times, the solar wind can cause the atmosphere's density to change, thus altering the drag and aerodynamic torques on the vehicle. These are factors which earthbound designers cannot predict and even after three decades of space flight, attitude control is still a major problem.

We were asked to make a trial of rule-based control, using a computer simulation of an orbiting space-craft under "black box" conditions. By this we mean that knowledge of the simulation's structure and parameters was unavailable to us and hence to the controller. Constraints and assumptions built into the client's specification of requirement included minimal human supervision. Only one ground station is to be used for control. The ground crew therefore have only a 16-minute window in each 90-minute orbit during which they can communicate with the space-craft. A premium is thus placed on the controller's aptness for generating intelligible reports.

The client's "black box" simulated three-axis rigid body attitude control with three non-linear coupled second order differential equations, and was supplied as Fortran object code. The use of pseudo-random generators introduced various time-varying disturbances, not only concerned with aerodynamic effects of solar wind variations and of atmospheric density and altitude changes, but also effects of propellant expenditure, payload redistribution, solar array articulation, extension and retraction of the gravity gradient boom and the motion of robotic and other on-board manufacturing appliances. Due to such unpredictabilities and to the possibility of a failure while out of communication with the ground, interest in a rule-based back-up controller centered on robustness, simplicity and conceptual transparency.

We already enjoyed familiarity with the BOXES adaptive rule-based control algorithm (Michie and Chambers, 1968; Chambers and Michie 1969). A laboratory problem of learning to balance a pole

was recently the subject of new work by Sammut (1988) who also reviewed trials of other algorithms for learning rule-based solutions to the problem. A rigid pole is hinged to a cart which is free to move along a track of fixed length. The learning system attempts to keep the pole balanced, and the cart within the limits of the track, by applying to the cart a force of constant magnitude but variable sign, either right or left ("bang-bang" control). The pole and cart system is characterised by four state variables which make up a four-dimensional space. By dividing each dimension into intervals, the state space is filled by four-dimensional "boxes". With each box (i.e. local region of state-space, or "situation" in the terminology of situation-action rules) is associated a setting which indicates that for any point within the given box the cart should be pushed either to the left or to the right. We set out to test essentially this representation on the client's simulated spacecraft.

The Black Box

The task was to drive the system from its initial state to the specified final state and maintain that state. Included in the black box is a fourth order Runge-Kutta numerical algorithm which integrates the dynamics of the equations of motion. The time step has a fixed value of 10 seconds. The black box keeps track of time and randomly injects various time-dependent disturbances as earlier described.

The state variables:

Attitudes: yaw (x), roll (y), pitch (z)

Body rates: ω_x , ω_y , ω_z .

Initial values of the state variables:

$x = y = z = 10$ deg

$\omega_x = \omega_y = \omega_z = 0.025$ deg/sec

The desired state:

$x = y = z = 0 \pm 3$ deg

$\omega_x = \omega_y = \omega_z = 0.005$ deg/sec

Failure conditions:

x or y or z exceeds ± 30 deg

ω_x or ω_y or ω_z exceeds ± 0.05 deg/sec

A flag is turned on if any of these go out of bounds.

Available control inputs:

Torque: T_x , T_y , T_z .

Torque is applied by the firing of thrusters which are aligned to the body axes. Although other attitude control devices (momentum exchange systems) will be used on the satellite in addition to thrusters, this work only addressed the use of thrusters. The following are minimum and maximum torques which can be applied by the thrusters:

$$T_x(\text{Min}) = T_y(\text{Min}) = T_z(\text{Min}) = 0 \text{ ft-lbf}$$

$$T_x(\text{Max}) = \pm 0.5 \text{ ft/lbf}; T_y(\text{Max}) = T_z(\text{Max}) = \pm 1.5 \text{ ft/lbf.}$$

The Rules

The first trial was made by directly adapting a set of BOXES-derived rules from the pole-and-cart domain to a sequential logic suggested by hand-derived rules due to Makarovic (1987, 1991). In each recognise-act cycle rule-matching follows a certain priority order, cycling through the state variables

until an action is selected. For each in turn the rule first checks that the first derivative does not exceed certain bounds. If it does, then a force is applied to oppose it. If it does not, then with respect to the same variable check its magnitude. If it exceeds given bounds then a force is applied accordingly.

In the case of the pole and cart, there was a clear priority to the order in which dimensions were checked. It was critical that the angular velocity and the angle of the pole were considered before the cart variables, since neglect of the pole leads to failure much more rapidly than neglecting to keep the cart away from the ends of the track. If this principle is applicable to the case of the space-craft then it is necessary to determine which of the state variables changes most rapidly. This was done, yielding rules expressible in "if-then-else" form, thus:

if $\omega_z < -0.002$ then	apply a T_z of +1.5
else if $\omega_z > 0.002$ then	apply a T_z of -1.5
else if $z < -2$ then	apply a T_z of +1.5
else if $z > 2$ then	apply a T_z of -1.5
else if $\omega_y < -0.002$ then	apply a T_y of +1.5
. . . and so on .	

Note that we are using "bang-bang" control, i.e. the torquers are either fully positive or fully negative just as in the pole-balancing experiments. With the space vehicle there are three dimensions, not one, to which a control motor (torquer) can apply a positive or negative thrust, corresponding to the yaw, roll and pitch dimensions of rotation respectively. The thresholds for the variables were determined by choosing an arbitrary value slightly within the bounds given for the desired values of the variables.

This control strategy proved to be successful but slow, requiring 8,700 seconds to bring the vehicle within desired bounds, and it also consumed 11.2 units of propellant. The question arose whether the control of each dimension could be decoupled. The cited rule only allows one thruster to be fired at any one time. If each axis of the craft were considered separately then all three thrusters could be fired simultaneously. This modification resulted in rules which brought the vehicle under control very quickly, requiring only 4,090 seconds. But propellant consumption, although improved, was still too high, using 7.68 units before the vehicle became stable. Therefore we decided on a partial retreat from pure "bang-bang", with a view to replacing it with finer control of the thrusters.

The resulting strategy is best understood by a decision array. For example, yaw control can be displayed as in Table 2 and the resulting performance as in Figure 1. Each of the 15 boxes corresponds to one control rule. Thus the box in the top left hand corner states that if the yaw is positive (i.e. above the bounds on the desirable yaw) and the yaw rate ω_x is well below the bounds of desirability then apply a quarter of the full torque in the positive direction. Thresholds were set for angles at ± 2 deg and for angular velocities they were ± 0.002 and ± 0.003 . The decision arrays for roll and pitch dimensions were of the same form. The resulting control behaviour was highly satisfactory. The pitch dimension was slowest of the three to be brought within the desirability zone.

Yaw too-positive	$T_x/4$	0	$-T_x/4$	$-T_x/2$	$-T_x$
Yaw OK	$T_x/2$	0	0	0	$-T_x/2$
Yaw too-negative	T_x	$T_x/2$	$T_x/4$	0	$-T_x/4$
	Yaw-rate too-neg.	Yaw-rate negative	Yaw-rate OK	Yaw-rate positive	Yaw-rate too-pos.

Table 2. A decision array for control of the yaw dimension.




Figure 1. Plot over time of vehicle's yaw behaviour(see text).


The client's engineers informed us that both in speed of recovery and in propellant expenditure our results were close to calculated optima. Since however it appeared that the satellite had greater inertia in the z -axis (pitch) than in the other two we decided to increase somewhat the thrust of the z -torquer. This brought the vehicle under control in 5,290 seconds, somewhat more slowly than the previous controller. But it only required 1.7 units of propellant, a substantial saving. Also calculations and simulations by the client's engineers made our result appear slightly better than optimal. This doubtless arose from minor approximations and/or distributional assumptions made in their numerical work. Time did not permit the point to be elucidated. But the broad conclusion was seen as extremely encouraging. An industrial-strength problem had shown that the simplicity, robustness and conceptual transparency of rule-based control does not have to be purchased at the cost of significant degradation of performance.

EXPERIMENTS WITH SKILL-GRAFTING

Supported by the freedom interactively to test each conjectured modification on the simulator, Sammut and I found our own powers of inductive conjecture adequate. But tasks of higher complexity, such as remote control of pilotless aircraft, would demand a less primitive approach. Present ideas are oriented towards the industry's use of interactive simulators for training pilots. A simulator-trained performer cannot tell you his or her strategy, but can demonstrate it. What is demonstrated can be automatically recorded. What is recorded can be inductively analysed by computer. With psychology-trained colleagues, Michael Bain, Jean Hayes-Michie and Chris Robertson I engaged in an investigation into the use of the rule-induction algorithm C4.5 (see Quinlan, 1987) to uncover effective control rules from such behavioural records. Experimental subjects were trained on an interactive simulation of a task illustrated in Figure 2. Control was exercised through a joystick of a pole-and-cart simulation which refreshed the screen approximately 20 times per second. New results together with earlier findings with this experimental system (Chambers and Michie 1969) lead to conclusions as follows (details are available in Michie, Bain and Hayes-Michie, 1990).

Conclusions from Pole-balancing

First conclusion: role of problem representation. Chambers and Michie used two regimes of training, identical except for the graphical animation seen by the subject. In one variant the picture was as shown. In the other the subject saw only a display of four separate horizontal lines, along each of which a pointer wandered to and fro. The subjects in this second variant were kept in ignorance of the nature of the simulated physical system. Unknown to them, the pointers actually represented the current status of four state variables, namely position of cart, velocity of cart, angle of pole and angular velocity of pole. Our hypothesis was that when the system is run fast, leaving only time for use and up-dating of procedural memory, then there will be no difference in the learning curves of subjects using the two different representations. Although not explicitly reported in their paper, an indication of this was observed by Chambers and Michie. In recent work we additionally used a rate sufficiently slow for subjects to report the task as having a major "planning" component. This slow-trained group learned more slowly, at least in the initial stages. In the new work trials have not yet been made of the lines-and-pointers representation.



Arrow indicates current direction of motor

Figure 2. Diagram of the pole and cart task.

Second conclusion: induction of rules from behaviour. Machine learning by imitation of a trained human was first shown for the inverted pendulum by Donaldson (1960) and reproduced under bang-bang conditions by Widrow and Smith (1964). Our concern was to test the ability of modern induction algorithms to extract from the behavioural record the kinds of rules believed to accumulate in procedural memory during skill-learning. Results have been positive. A task was investigated where the object was to cross the centre of the track as often as possible in an allotted time-span without dropping the pole or crashing the cart. When induction-extracted rules were installed in the computer as an “auto-pilot”, performance on the task was similar to that of the trained human who had generated the original behavioural trace, but more dependable, as described below.

Third conclusion: the clean-up effect. Rules induced from a behavioural record can be assessed in two different ways. Predictive mode tests the ability of a rule-set correctly to predict other behaviour sampled from the same source. Performance mode tests the ability of the rules to substitute for the human source in executing the skilled task.

Induced rule-sets performed satisfactorily in the second mode while consistently showing high error rates, often exceeding 20%, in predictive mode. My co-worker Mr. Michael Bain pointed out that when watching a machine-generated rule-set’s performance on the screen one is struck by an appearance of superhuman precision and stability. We surmised that a trained human skill, although controlled by an equally precise and stable set of production rules, is obliged to execute *via* an error-prone sensorimotor system. The superimposition of inconsistencies and lapses of attention would then be stripped away by the averaging effect implicit in inductive generalization, thus restoring to the experimenters a cleaned-up version of the original production rules. When tested in predictive mode, such a rule-set can do no better than the cumulative sum of human perceptual and execution errors allow. But in performance mode one would expect a super-reliable stereotype of the behaviour of the human exemplar. Direct confirmation of this idea was obtained by calculating the magnitude of the pole and cart’s excursions during a control session along each of the four dimensions of the state space. Observed ranges tabulated in Table 3 were obtained from a behavioural trace recorded from Mr. Bain’s own trained performance.

	x	x-dot	theta	theta-dot
Trained human (ranges)	2.79	4.85	0.562	5.021
Induced rule (ranges)	0.46	1.83	0.134	2.276
Range differences	2.33	3.02	0.428	2.745
“Clean-up”	83%	62%	76%	55%

Table 3. Clean-up effect shown by induced control rules over a 5-minute test period.

x = position, theta = angle: “dot” denotes first derivatives. These results are typical, and have been many times confirmed in test runs with the same, and with other, subjects.

The findings suggest that “skill-grafting” from behavioural traces may be possible for more demanding tasks, such as those encountered in aircraft flight control. The key idea is that if we could look inside the head of the ground-based pilot of a remotely controlled aircraft, or of the on-board pilot of a difficult vehicle such as a helicopter, we might see a neural encoding of a fully sufficient skill, but degraded in real-time execution by sensorimotor delays and errors. Recovery of a logically equivalent rule structure and its transplantation to an error-free device (i.e. to a control computer) then offers a source of enhanced and more reliable performance. In advanced rotorcraft control there is a current need for libraries of individual autopilot manoeuvres (“circle at 50 feet”, “fly slowly sideways for one minute”, etc.) which the pilot could activate in difficult weather or other conditions, so as to free his attention for some main task in hand, – visual search of water surface, target acquisition, etc.

Learning to Fly

Sammut and colleagues have recently been able to reproduce the “skill-grafting” phenomenon in the complex task of flying a simulated aircraft (Sammut, Hurst, Kedzier, and Michie, 1992). Using a flight simulator developed by Silicon Graphics, three subjects trained themselves by repeatedly piloting a simulated Cessna through the successive stages of a defined flight plan, consisting of the following manoeuvres:

1. Take off and fly to an altitude of 2000 feet.
2. Level out and fly to a distance of 32,000 feet from the starting point.
3. Turn right to a compass heading of approximately 330° .
4. At a North/South distance of 42,000 feet, turn left to head back towards the runway.
5. Line up on the runway.
6. Descend to the runway, keeping in line.
7. Land on the runway.

Taking ‘events’ as being signalled by the occurrence of control actions, then up to 1,000 events were recorded per flight. Each of three trained subjects performed 30 flights, so that the complete data comprised about 90,000 events. For each event the control action was recorded, together with values of state variables measured at a moment selected 1-3 seconds earlier. The “offset” makes approximate allowance for the pilot’s delay in responding to complex stimuli. To give a rough impression of the data, the following are names of recorded variables:

boolean variables: *on-ground, g-limit, wing-stall*;

integer variables: *twist, elevation, azimuth, roll-speed, elevation-speed, azimuth-speed, air-speed, climb-speed, fuel, thrust, flaps*;

real variables: *E/W distance, altitude, N/S distance, rollers, elevator*.

The simulation program was modified to log the subjects’ actions during flight. Log files from trained subjects were used to create the input to an inductive rule-learning program. The learning program’s output took the form of separate decision trees for each of the 4 different control actions, further sub-divided into the 7 stages listed above. For example, to quote from the original paper,

The critical rule at take-off is the elevator rule:

elevation > 4: level-pitch

elevation ≤ 4:

airspeed ≤ 0: level-pitch

airspeed > 0: pitch-up-5

This states that as thrust is applied and the elevation is level, pull back on the stick until the elevation increases to 4° . Because of the delay, the final elevation usually reaches 11° which is close to the values usually obtained by the pilot. "pitch-up-5" indicates a large elevator action, whereas "pitch-up-1" would indicate a gentle elevator action. The other significant control at this stage is flaps:

elevation ≤ 6 : full-flaps

elevation > 6 : no-flaps

Once the aircraft has reached an elevation angle of 6° , the flaps are raised.

The 28 decision trees were automatically converted to C-code routines, arranged as a suite of 7 flight control modules, each responsible for all aspects of a given stage. A new module was invoked as soon as a pre-programmed precondition was satisfied for the onset of the next stage. Within each module, 4 sets of if-then rules separately supervised the four separate control actions.

An autopilot was generated in this fashion from each of the trained subjects. Tests were made by running the simulator in autopilot mode, substituting as autopilot code one or another of the three inductively synthesized program suites. All three executed the entire flight plan with conspicuous competence, but with individual mannerisms characteristic of the flying styles of the respective human data sources. Indications of the "clean-up effect" (see earlier) were also evident, particularly during the approach stage.

Conclusions from Autopilot Induction

Models extracted as above from decision-data by rule learning are purely heuristic in form. They incorporate no explicit references to time or causality. Yet as reviewed earlier, real-time human problem-solving involves co-operation between two separate kinds of mental process driven from two separate memory systems, updated by separate kinds of learning. The dichotomy is recognised in AI under the labels "heuristic" and "causal". These roughly correspond to the neuropsychologist's procedural/declarative distinction. The balance in practice is set by the time-constraints imposed by different tasks. A fast situation-action cycle allows time only for executing heuristics and virtually none for reasoning about causes.

As the autopilot experiments demonstrate, complex skills *can* be built entirely from heuristics. Bears can learn to ride bicycles, — and humans can fly combat planes through mission phases which allow no time for analysis. Under such circumstances, everything goes by pattern-invocation. The formal identity between pattern classification and control then stands out clearly. This identity has recently been discussed (Michie, 1991) in connection with a definition of learning which says:

a learning system uses sample data (the training set) to generate an up-dated basis for improved classification of subsequent data from the same source.

The above-cited discussion continues:

Notice that the definition, although phrased strictly in terms of classification, logically extends to acquisition of improved performance on tasks which do not look at all like classification. Iterative situation-action tasks come to mind such as riding a bicycle, solving an equation, or parsing a sentence. The extension becomes obvious when for the decision classes we choose names which refer to

partitions of the space of situations as 'suitable for action A', 'suitable for action B', etc.

Why should one want, in addition to finding a machine-efficient representation of the above mapping, to construct an operationally redundant superstructure to capture causal relations and to support "what-if" planning? Answers suggest themselves as soon as one moves to the more demanding definition which animates the characteristically AI approach to learning:

a learning system uses sample data to generate an up-dated basis for improved classification of subsequent data from the same source *and expresses the new basis in intelligible symbolic form.*

The requirement for social communication of the "improved basis" now forces the issue. If synthetic autopilots are to show "understanding" of flight situations and their own responses, then however necessary heuristic models may continue to be for the sub-structures of skill, insightful performance and explanation at higher and more strategic levels demands causal modelling of a sophisticated kind. It is towards this difficult objective that much work of the kind here reviewed is now turning (see Bratko, 1992).

SUMMING UP

In the debate between symbolic and neural-net representations, the two sides have tended to overlook the possibility that different parts of the brain, specialized to address different purposes, employ different representations. Specifically such differing purposes can be broadly grouped under two contrasted main heads:

- (1) "run-time" thinking;
- (2) communication of the process and its outcome.

For (1), there is no obvious biological reason to expect symbolic representations to have evolved, in the sense in which "symbolic" is here used. Indeed there is little evidence that such structures are employed in the brain's real-time problem-solving, much of which is critically supported by varieties of visual and spatial reasoning associated with the brain's right cerebral cortex. But the *social* dissemination of knowledge and thought listed under (2) is of such predominant importance in our species that elaborate symbolic mechanisms have emerged to support the execution of this function. The evolutionary processes have partly been biological and partly cultural. Eccles (Popper and Eccles, 1977) paints a picture of divergent specialization between the two hemispheres of the brain according to which the "minor" (usually the right) hemisphere plays roles central to run-time problem-solving, involving pattern-handling and spatial and social orientation. Yet this hemisphere almost wholly lacks capabilities of symbolic reasoning, notably those associated with language and logic. The dominant (usually left) hemisphere, by contrast, not only fluently handles the decipherment of linguistic and logical expressions, but is also the clearing-house for reports on subgoal attainment during problem-solving. Eccles argues that, although "consciousness" is also manifested by the right hemisphere in the sense of a diffuse awareness, the focussed and organised forms of goal-oriented awareness which we associate with "self" are the sole prerogative of the left brain. More recently the possibility has been aired in neurobiological circles (see Benjamin Libet's observations and associated discussion in *Behavioral and Brain Science*, 1988-89) that the seat of consciousness acts more as a news room than as a planning headquarters, putting a coherent retrospective gloss on the consequences of decision. The decisions themselves, in this model, emanate from activities localised elsewhere. An elaboration of this view has recently been developed by Dennett (1992).

Whatever the neural nature of functions (1) and (2) above, modern brain science sees them as operationally and topographically distinct. In such a view, the mechanisms of (2) face a serious problem. Modules specialized to symbolic reporting must interface with dissimilar, even alien, architectures if explanations of the "self's" problem-solving decisions are to be generated. Little wonder, when required to do service in the more intuitive field of real-time skills, the brain's explanation module should be found wanting, — or should resort, when pressed by the dialogue-elicitation specialist, to confabulation.

Are we, as engineers of cognition, obliged to burden intelligent artifacts with similar problems? On the contrary, to do so would seem the height of folly. Moreover, from such work as has here been reviewed, an alternative strategy is available. We can treat expert subcognition as a "black box" from which articulate models can be extracted. The product: symbolic models of sub-symbolic behaviour, or, more concretely, machine-executable yet articulate skills from "silent" brains.

ACKNOWLEDGEMENTS

Thanks are due to Professor Quinlan for active assistance in the use of his C4.5 algorithm. The preparation of the review also benefited from criticisms and suggestions made by Dr. Patricia Smith Churchland, Mrs. J.E. Hayes Michie, Prof Colwyn Trevarthen, and by colleagues in the helicopter division of the Royal Aerospace Establishment, Bedford, UK, and at Advanced Rotorcraft Technology Inc., Mountain View, USA. I was helped by facilities at the Turing Institute, UK, and at the Department of Computer Science, University of New South Wales, Australia. The final draft was completed during my tenure of the C.C. Garvin Endowed Visiting Professorship at Virginia Tech, USA. Thanks are also due to the Japan Society for Artificial Intelligence, in whose 1991 Proceedings some of the material of this article first appeared.

REFERENCES

- Anderson, J.R. (1990). *Cognitive Psychology and its Implications*, (Third Edition) W.H. Freeman & Co.
- Bratko, I. (1991). Qualitative modelling: learning and control. *Proc. AI-91*, Prague. Copies also available from Electr. Eng. and Comp. Sci., Ljubljana University, Slovenia.
- Chambers, R.A. and Michie, D. (1969). Man-machine co-operation on a learning task. In *Computer Graphics: Techniques and Applications* (ed. R. Parslow, R. Prowse and R. Elliott-Green), London: Plenum.
- Dennett, D.C. (1992). *Consciousness Explained*. Little, Brown and Co.
- Donaldson, P.E.K. (1960). Error decorrelation: a technique for matching a class of functions, in *Proc. Third Internat. Conf. on Medic. Electronics*, 173-178.
- Eccles, J.C. (1977). In *The Self and its Brain* (authors and editors: K.R. Popper and J.C. Eccles), London and New York: Routledge and Kegan Paul.
- Feigenbaum, E.A. and McCorduck, P. (1983). *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Reading, MA: Addison-Wesley.
- French, R.M. (in press). Subcognition and the limits of the Turing Test. In *Mind*.
- Ginsberg, M.L., (ed, 1987). *Readings in Nonmonotonic Reasoning*, Los Altos, CA: Morgan Kaufman.
- McCarthy, J. (1959). Programs with common sense. In *Mechanization of Thought Processes* Vol. I, London: Her Majesty's Stationary Office. Reprinted in M. Minsky, ed. (1960) *Semantic Information Processing*, Cambridge, MA: MIT Press.
- Makarovic, A. (1987). Pole-balancing as a benchmark problem for qualitative modelling. *Technical Report DP-4953*, Ljubljana: Josel Stefan Institute. Revised as (1991): A qualitative way of solving the pole-balancing problem. In *Machine Intelligence 12* (eds. J.E. Hayes, D. Michie and E. Tyugu), Oxford University Press.
- Michie, D. (1988). The Fifth Generation's unbridged gap. In *A Half-Century of the Universal Turing Machine* (ed. R. Herken), Oxford University Press.
- Michie, D. (1989). Problems of automatic concept formation. In *Applications of Expert Systems*, Vol.1 (ed J.R. Quinlan), Addison-Wesley, pp.310-333.
- Michie, D. (1991). Methodologies from machine learning in data analysis and software. *Computer Journal*, **34**, 559-565.
- Michie, D., Bain, M. and Hayes-Michie, J. E. (1990). Cognitive models from subcognitive skills. In *Knowledge-based Systems in Industrial Control* (ed. Grimble, M., McGhee, S. and Mowforth, P.), Peter Peregrinus.

- Michie, D. and Chambers, R.A. (1968) BOXES: an experiment in adaptive control, In *Machine Intelligence 2* (eds. E. Dale and D. Michie), Edinburgh: Edinburgh University Press.
- Newell, A. and Simon, H.A. (1976). Computer science as empirical inquiry: symbols and search. *Commun. of the ACM*, **19**, 113-126.
- Popper, K.R. and Eccles, J.C. (1977). *The Self and its Brain*. London and New York: Routledge and Kegan Paul.
- Posner, M.I. (1973). *Cognition: an Introduction*, Glenview, IL: Scott, Foresman.
- Quinlan, J.R. (1987). Generating production rules from decision trees. *Internat. Joint Conf. on Art. Intell. 1987 (IJCAI-87)*, Los Altos, CA: Kaufmann, pp.304-307.
- Sammur, C. (1988). Experimental results from an evaluation of algorithms that learn to control dynamic systems. In *Proc. Fifth Internat. Conf. on Machine Learning* (ed. J. Laird), San Mateo: Morgan Kaufmann.
- Sammur, C., Hurst, S., Kedzier, D. and Michie, D. (1992). Learning to fly. In *Proc. Ninth Intern. Machine Learning Conf.* (ed. D.H. Sleeman).
- Sammur, C. and Michie, D. (1991). Controlling a "black box" simulation of a space craft. *The AI Magazine*, Amer. Assoc. for Artif. Intell.
- Simon, H.A (in press) Machine as mind, In *Proceedings of the Turing 1990 Colloquium*, 3-6 April 1990, Brighton, UK, (ed Millican, P.) to appear.
- Slocombe, S., Moore, K. and Zelouf, M. (1986). Engineering expert system applications. Presented at *BCS Annual Conference*, December 1986.
- Squire, L. (1987). *Memory and Brain*, Oxford University Press.
- Sterling, L. and Shapiro, E. (1986). *The Art of Prolog*, Cambridge, MA: The MIT Press, p.357.
- Widrow, B. and Smith, F.W (1964). Pattern recognising control systems, in *Computer and Information Sciences* (eds Tou, J.T. and Wilcox, R.H.). Clever Hume Press.