This article was downloaded by: [University of Calgary]
On: 14 February 2012, At: 10:52
Publisher: Taylor & Francis
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK

International Journal of General Systems
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/ggen20

SYSTEM IDENTIFICATION, APPROXIMATION AND COMPLEXITY

BRIAN R. GAINES [a]

[a] Man-Machine Systems Laboratory, Department of Electrical Engineering, University, of Essex, Colchester, Essex, U.K

Available online: 06 Apr 2007

To cite this article: BRIAN R. GAINES (1977): SYSTEM IDENTIFICATION, APPROXIMATION AND COMPLEXITY, International Journal of General Systems, 3:3, 145-174

To link to this article: http://dx.doi.org/10.1080/03081077708934759

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: http://www.tandfonline.com/page/terms-and-conditions

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Int. J. General Systems
1977, Vol. 3, pp. 145–174

© Gordon and Breach Science Publishers Ltd., 1977
Printed in Great Britain

SYSTEM IDENTIFICATION, APPROXIMATION AND COMPLEXITY

BRIAN R. GAINES

Man-Machine Systems Laboratory, Department of Electrical Engineering, University, of Essex, Colchester, Essex, U.K.

(*Received March 9, 1976; in final form April 12, 1976*)

This paper is concerned with establishing broadly based system theoretic foundations and practical techniques for the problem of system identification that are rigorous, intuitively clear and conceptually powerful. A general formulation is first given in which two order relations are postulated on a class of models: a constant one of *complexity* and a variable one of *approximation* induced by an observed behaviour. An *admissible model* is such that any less complex model is a worse approximation. The general problem of identification is that of finding the *admissible subspace* of models induced by a given behaviour. It is proved under very general assumptions that, if deterministic models are required, then nearly all behaviours require models of nearly maximum complexity. A general theory of approximation between models and behaviour is then developed based on subjective probability concepts and semantic information theory. The role of structural constraints such as causality, locality, finite memory, etc., are then discussed as *rules of the game*. These concepts and results are applied to the specific problem of stochastic automaton, or grammar, inference. Computational results are given to demonstrate that the theory is complete and fully operational. Finally the formulation of identification proposed in this paper is analysed in terms of Klir's epistemological hierarchy and both are discussed in terms of the rich philosophical literature on the acquisition of knowledge.

INDEX TERMS Systems identification, systems approximation, systems complexity, admissible model, epistemological hierarchy, grammatical inference, logical probability.

Downloaded by [University of Calgary] at 10:52 14 February 2012

1 INTRODUCTION

The problem of inferring the structure of a system from observations of its behaviour is an ancient one with many ramifications. The literature on the subject is vast, having its roots in the philosophical problems of the nature of *reality* and our inference of it from sensations.[1,2,3] Plato's famous simile[4] of the prisoners in the cave who, "like ourselves... see only their shadows, or the shadows of one another, which the fire throws on the opposite wall of the cave", epitomizes the inherent uncertainty and tenuous nature of inference processes that we all perform and take for granted in our everyday life.

Even if the metaphysical problems of *reality* are discarded for a more pragmatic approach that asks, not whether our inferred structures are *real*, but instead whether they *work*, i.e. are useful, or valid in some weaker sense, deep philosophical problems remain. If our model structure accounts only for the observed behaviour on which it is based then it appears to suffer from the usual defect of deductive inference, that is inferentially vacuous leaving us with only a re-description of the observations. Whereas if we demand that our structures be *predictive*, allowing us to account for further observations, then we come up against Hume's[5–6] conclusive arguments that inductive inference can never be validated (at least, neither deductively nor inductively).

Much of the literature on the "philosophy of science" is concerned with the epistemological problems stemming from Hume's arguments. Reactions range from existentialist dread,[7] emphasizing the personal and unique nature of our experience and evaluation of it, to the totally impersonal formal methodologies of incremental date acquisition and model structure adjustment of Carnap's logical probability and confirmation theory.[8–13]

Within this spectrum one has a range of existentialist positions;[14] the deep methodological studies of Brentano[15] and Husserl[16] leading to a variety of *phenomenological* analyses;[17] arguments that observational descriptions of phenomena are already *theory-laden*;[18–21] logical analyses of the role of *analogy* in model formation,[22–24] of *simplicity* and *economy* in model selection,[25–27] and of *convention* in model utility;[28] the demonstrations that *de-*

duction is itself open to Hume's criticisms;[29–30] the detailed analysis of *specific flaws* in the texture of his argument;[31–32] the aphoristic *metaphysical reply* to Hume given by Wittgenstein;[33–34] the various *vindications* of induction developed by Reichenbach,[35] Harrod,[36] Katz,[37] Black,[38] Rescher,[39] and others;[40] Popper's[41–44] methodological point that hypothesized structures can never be verified, only *falsified*, and the Popper-Carnap controversy[45] over falsification versus confirmation; the sociological models of scientific *revolutions* of Kuhn[46] and the painstaking studies of how these social functions actually operate of Merton;[47] the more *structural* rationale of scientific method of Lakatos[48] and Hesse;[49] the *anarchistic* counter examples of Feyerabend,[50] and the cooly cynical appraisal of the *arbitrariness* of the whole debate by Gellner.[51]

Many aspects of these philosophical debates find a more mathematical formulation in studies of the history[52] and foundations of *probability theory*,[12,35,53–59] *semantic information theory*,[60–63] *decision making and statistical inference*,[64–66] *computational complexity*[67–72] and *grammatical inference*.[73–76] The philosophical debate has also become rather more pointed in recent years because it has been possible to apply the arguments *operationally* to machines[77] that exhibit reasoning by analogy,[78] law discovery[79,80] inferencing from incremental observations,[81,82] expectation,[83] and many other manifestations of human learning behaviour.[84]

With this background in mind one treads warily in the development of behaviour-structure inferencing techniques in system theory. Many of the philosophical problems may be evaded by assuming that the observed behaviour arose from one of a set of known possible systems. For example, Zadeh[85] defines *system identification* as "the determination on the basis of input and output, of a system within a specified class of systems, to which the system under test is equivalent". Most practical studies[86] of system identification operate within this framework and presuppose both a "specified class of systems" and a well-defined decision procedure for determining that a system is "equivalent" to one of these systems on the basis of its input-output behaviour.

I shall adopt this point of view in this paper and give a general systems theoretic formulation of the identification problem that encompasses previous specific formulations and yet is sufficiently precise for some interesting features of identification to be determined. In particular the formulation deals very effectively with a major problem left open in the definition above—that of defining "equivalence" when the class of systems considered is acausal in some sense, e.g. nondeterministic or stochastic automata. The input-output behaviour of an acausal system is *not* uniquely related to its structure and some element of approximation[87] is essential in the definition of equivalence. The formulation proposed allows for this in a very clear and general form that enables, for example, the problem of identifying stochastic automata, or grammars, to be rigorously defined and solved.[88–90]

Out of the general formulation come two specific lines of development:—

a) How can the techniques be applied to specific classes of system and what theoretical and practical results may be obtained, i.e. given that the source of behaviour is a deterministic finite state automaton (DFSA) or a stochastic FSA (SFSA) what are an appropriate definition of equivalence, and a practical implementation, and what results may be obtained about them, theoretically and experimentally?

b) How should the class of model systems be specified and what happens if the observed behaviour arises from a system *not* in this class? e.g. will the identification system behave reasonably and give some approximate and useful results, or does it fail completely. Here there are some surprising results that run counter to our (current) intuitions.

Both aspects of the identification problem are developed in this paper. In particular, the problem of identifying stochastic finite state automata and grammars is analysed theoretically and practical experimental results are given. These are closely related to the literature on computational complexity already noted, but have additional interest because they also establish close links with *subjective* foundations of probability,[53,58,91–100] in particular Pearl's recent results on economic bases for subjective probability[100] and relations between approximation and complexity.[101–103] Also the results of using classes of model systems different from the system observed are analysed and discussed for a variety of cases. The various aspects of the problem are also linked to the rich philosophical literature on the acquisition of knowledge referenced above. This is made more explicit and precise than was previously possible by using the hierarchical classification of epistemological levels recently proposed by Klir[104] as a framework for both the philosophical connotations and the identification techniques.

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

This paper is one of a series[88–90] concerned with establishing broadly based system-theoretic foundations and practical techniques for the problem of system identification that are rigorous, intuitively clear and conceptually powerful. An informal introduction to the techniques and a range of experimental examples have been given in Ref. 89, and studies have been reported of applications to the inference of finite-state probabilistic grammars.[90] It is expected that later papers will report particular applications to the analysis of sequential behaviour in humans and animals. The role of this paper is to provide the background and foundations for such studies.

2 A GENERAL FORMULATION OF THE IDENTIFICATION PROBLEM

Our problem can be stated initially to be: "given a sample of the behaviour of some system, to determine the optimum model from some prescribed models that would account for it". Note that Zadeh's definition of the previous section has already been generalized. We do not need at this stage to define what is meant by *behaviour*, the terms of reference by which we describe it, etc. In particular, the notions of *input* and *output* have been dropped. These are structural concepts that belong to the models *not* the behaviour, and in the examples given it will be shown that the input/output distinction can be *inferred* and need not be postulated. A similar remark has been made by Klir (Ref. 105, p. 166) who calls such systems without the input/output distinction *neutral*.

In addition the notion of *equivalence* between behaviour and structure has been dropped. This, in the sense of a mathematical equivalence relation, is too powerful a notion for a theory that must encompass the modelling of acausal systems. We need instead some concept of *degree of approximation*, an order relation that allows us to say that one model accounts for the behaviour *better* than does another. When the more powerful equivalence *does* exist, for example, in modelling the behaviour of deterministic finite-state sources with DFSA then it allows for the elegant category-theoretic formulation of identification in terms of an *adjunction* between behaviour and structure developed by Goguen,[106,107] Arbib,[108,109] and Ehrig.[110,111]

However, determinism is a myth in the real world, although the success of mechanistic models of the universe in the eighteenth century has made it something of a holy grail for science until the present day and there is substantial evidence that the assumption of deterministic causality is deeply rooted in both human cognitive and perceptual processes.[112] Many authors have argued for the replacement of deterministic causality with *probabilistic causality*.[113–116] I have proposed elsewhere[117] that the process in science that Carnap[8] calls *precisiation*, of replacing a phenomenal *explicandum* with a noumenal *explicatum* need not be interpreted in the narrow sense of precise deterministic explicata, and have demonstrated[118] that this is not only a metaphysical point but also a practical one—*the universe becomes incredibly complex and our models of it nonsensical if we assume determinism in the face of even a slight trace of acausal behaviour.*

2.1 *Complexity and Admissibility*

Having weakened *equivalence* to an order-relation of *approximation*, we face one residual problem with the definition above, that the concept of an *optimum* model is no longer appropriate. For example, imagine that you and I are each given the same sample of behaviour and asked to model if from the same class of models. "My model is a better approximation," I say, "Ah," you reply, "but mine is a far *simpler* model. Indeed, I am not sure that all yours does is not just to retain a *memory* of the behaviour, whereas mine, whilst a worse approximation, is clearly a far better representation. If the behaviour were actually being generated by the system corresponding to my model the degree of approximation I have achieved would be quite reasonable."

These are the key issues, that we do not rate all models as themselves equivalent. There is invariably an order relation on our prescribed class of models that gives rise to a trade-off between the degree of approximation and the preference for models. It is common to call this ordering one of *complexity* with the preference being for the less complex models. For convenience I shall adopt this terminology, but with the warning that *the order relation is not intrinsic to the class of models.* We may both adopt the same class of models but what I regard as complex may for you be simple. The ordering of models in terms of complexity is *arbitrary* and depends upon our individual points of view.

Too much stress cannot be laid upon the fact that our model classes are incompletely specified for purposes of identification until we have defined an

Downloaded by [University of Calgary] at 10:52 14 February 2012

ordering upon them. It is a trap into which we may easily fall, particularly in general systems theory. Several classes of models open to us appear so general that we feel they must be adequate to account for *all* possible behaviours. And so they are, but that is not sufficient to allow us to suppose that we can base an all-embracing system science upon this class of systems. When we specify our order relation upon the models we may find that the behaviours of many important systems require complex models under our ordering, whereas, with a different ordering on the same class of models, they all become simple.

If this happens then it is probable that there will be a *scientific revolution* in which the order relation on our models is changed to make the observed world less complex. Since we normally wish to associate the ordering with some intrinsic feature of our models this will also lead to us viewing the models in a different way so as to emphasize some new aspect of their structure. For example, clearly every finite sample of behaviour (which is all we ever have) can be accounted for by a DFSA. However, so can it by an SFSA or a Turing machine—why choose one rather than another? The results of Ref. 118 show that the behaviour of a simple stochastic FSA requires, in general, a very complex deterministic FSA model, and, likewise, the behaviour of a simple push-down automaton requires very complex models with both DFSA and SFSA.

Given order relations of: *complexity* on our class of models, and of *approximation* on the extent to which a model accounts for a given observed behaviour, it is possible to give a precise formulation of identification in a general systems context: "given a sample of the behaviour of some system, to determine those models from some prescribed class of models that are *admissible* in that, for each one, any other model that gives a better approximation in accounting for the behaviour is more complex." The concept of admissibility is one borrowed from statistics[119] and proves a powerful one in problems of control theory[120] and pattern recognition[121] in situations where no definition of optimality is possible. Note that there is rarely a unique admissible model but instead a subset containing several admissible models. However, this subset has some interesting properties that do much to make up for the lack of a single unique model.

In the next subsection I shall put this definition of identification into mathematical form, go on to analyse some of its properties, particularly those related to computational complexity and then develop more specific features of it related to particular classes of identification problems.

2.2 *Mathematical Formulation of Identification*

The concepts developed in the previous sections may be formulated more formally in terms of: a set of possible observed behaviour, B; a set of models, M; the pointed monoid, $(\mathbf{Ord}_M, \leqq)$, of all order relations on M with one specified relation, $\leqq$, singled out; and a mapping, $f: B \rightarrow \mathbf{Ord}_M$, from the set of behaviours, B, to the set or order relations M, $\mathbf{Ord}_M$. The quadruple $(B, M, \leqq, f)$, defines an *identification space*: the relation $\leqq$ is one of model complexity and if $m, n \in M$ are such that $m \leqq n$ we shall say that the model m is not more complex than n. Other considerations being equal it will be assumed that the least complex possible model is preferred. Note, however, that $\leqq$ may be only a partial order so that, in general, there will be a set of minimal models rather than a unique minimum. The mapping f is determined by the further order relation of approximation that each behaviour induces on the set of models. We shall write for $b \in B$, $\leqq_b = f(b) \in \mathbf{Ord}_M$, and if m, $n \in M$ are such that $m \leqq_b n$ we shall say that model m is not a worse approximation to behaviour b than is model n. The best models for b are thus those minimal in the order relation $\leqq_b$ which again need not be more than a partial order.

Now we are in a position to define a solution of the identification problem in terms of the product of the two order relations, $\leqq_b^*$:

$$\forall m, n \in M, m \leqq_b^* n \Leftrightarrow m \leqq n \text{ and } m \leqq_b n, \qquad (1)$$

i.e. $m \leqq_b^* n$ if and only if m is neither more complex nor a worse approximation than n. The minimal elements of the new order relation have the property that there are no other models that are both less complex and a better approximation than them. Even if both $\leqq$ and $\leqq_b$ are total orders it is likely that $\leqq_b^*$ will be a partial order (we can trade more complexity for better approximation) and hence there will be in general no unique minimum model. The minimal elements are all *admissible*[119] solutions to the identification problem because they cannot be decreased in complexity without worsening the approximation and cannot be improved in approximation without increasing complexity. Thus we may define the solution of the identification problem for a space $(B, M, \leqq, f)$ and an observed

Downloaded by [University of Calgary] at 10:52 14 February 2012

sequence $b \in B$ to be the *admissible subspace* determined by b, $M_b \subset M$, such that:

$$M_b \equiv \{m : \forall n \in M, n \leqq_b^* m \Rightarrow m \leqq_b^* n\} \qquad (2)$$

i.e. if any model is better than one in M_b then it is equivalent to it under the order relation, $\leqq_b^*$, the usual requirement for minimality. Thus models in M_b are either equivalent or incomparable in terms of $\leqq_b^*$. Note that, for rigour when dealing with infinite sets we should impose the constraint that the minimal elements under $\leqq_b^*$ exist and belong to M, e.g. by taking a suitable closure.

M_b may be shown to have similar, but inverse, structures under $\leqq$ and $\leqq_b$:

Result 1. For comparable models, the relations $\leqq$ and $\leqq_b$ are antitone in M_b, i.e. $\forall m, n \in M_b : m$ and n are comparable under $\leqq$ and $\leqq_b$, $m \leqq n \Leftrightarrow n \leqq_b m$.

Proof. Assume $m \leqq n$—if $m \leqq_b n$ then $m \leqq_b^* n$ so that, since $n \in M_b, n \leqq_b^* m$ which implies $n \leqq_b m$. Conversely, assume $m \leqq_b n$—if $m \leqq n$ then $m \leqq_b^* n$ so that, since $n \in M_b, n \leqq_b^* m$ which implies $n \leqq m$. Note the symmetric way in which the relations of complexity and approximation have entered the discussion and definitions, and the anti-symmetry between them in the admissible subspace of models shown by this result. Models in the admissible subspace are effectively unordered by $\leqq_b^*$ but ordered in effectively identical, but inverse, ways by $\leqq$ and $\leqq_b$.

2.2.1. *Example—identification of deterministic automata* It is these well structured subspaces that replace the unique solutions of, for example, problems amenable to the Nerode equivalence,[122,123] and it is of interest to see how such problems fit within the current framework. Consider the identification space $(D^*, M, \leqq, f)$, where D^* is the free monoid generated by a (finite) alphabet, D, of atomic *descriptors* (inputs or outputs); M is the set of irreducible, Mealy, finite state deterministic automata with a specified initial state having inputs and outputs in D; the ordering of complexity, $\leqq$, is determined by the number of states of the automata, so that $m \leqq n \Leftrightarrow$number of states of m less than or equal to the number of states of n; and the ordering, $\leqq_b$, induced by a sequence of observed input-output behaviour, $b \in D^*$, is in fact a binary classification in which m is maximal in the order unless it generates the sequence b exactly when it is minimal (where m starts in the specified initial state, receives the input sequence imbedded in b and emits an output sequence that is tested against that imbedded in b).

This is the standard case of deterministic automata inferencing from a sample of the input-output behaviour and the Nerode equivalence, or one of its generalizations, may be used to determine a minimal-state machine[124] using high-speed algorithms.[125] This essentially splits the space of possible machines into three sets:

1) a unique machine (up to isomorphism) that is minimal in the ordering of approximation (an exact fit) and, subject to this, minimal in the ordering of complexity (minimal state).

2) A set of machines with fewer states than this that are maximal in the ordering of approximation, i.e. are simpler but do not fit the behaviour.

3) the remaining machines with more states, or with the same number of states that do not fit the behaviour.

The first two sets of machines together form the admissible subspace, M_b, for the problem. In this example the second set of machines is of little interest because there is no gradation of approximation. It would be possible to define a graded form of approximation in terms of some finer evaluation of the extent to which the outputs of a model, m, diverge from those of the behaviour, b. However, in conventional deterministic modelling it is the uniqueness of the "solution" obtained in the first set of admissible models that is of interest.

2.2.2. *Example—identification of probabilistic automata* The problem of identifying probabilistic automata will be treated in more detail in sections 3.3.1 and 4. However, it is useful to contrast the techniques I have previously described [88,89] with those for deterministic automata above. The main difference between the two problems is that a probabilistic automaton model gives not a specific output but instead a *distribution over possible outputs*, and a distribution over possible next states. We can evaluate the distribution over outputs with respect to the actual observed output by using one of the loss functions devized to elicit subjective probabilities,[93–100] e.g. a loss of minus the log of the proposed probability for the output that actually occurs (this is zero if the actual output is predicted with probability 1 and positive otherwise). We can eliminate the effect of having only a distribution over next states by using *observable* automata only in which the actual output that occurs is sufficient to resolve the uncertainty as to the next state.

Downloaded by [University of Calgary] at 10:52 14 February 2012

Consider the identification space, $(D^*, M, \leqq, f)$, where: D^* is a free monoid as before; M is now the set of irreducible, observable, Mealy, finite-state probabilistic automata over D with a specified initial state; the ordering of complexity, $\leqq$, is number of states as before; and the ordering, $\leqq_b$, induced by a sequence of observed input-output behaviour, b, is determined by the natural numerical ordering on the sum of the losses when a model, m, is used to predict b (where m starts in the specified state, receives the input sequence imbedded in b and emits probability distribution over the outputs that are used in conjunction with the actual output to determine both the loss and the next state). The smallest loss gives the best approximation and zero loss (minimum possible) corresponds to exact deterministic prediction of the outputs and hence to the deterministic modelling already discussed.

The admissible subspace for probabilistic identification does not split trivially as it did for deterministic modelling. For a given number of states there will generally be models that give a smaller loss (better approximation) than any models with fewer states. As the number of states in the model (the complexity) increases the loss will get less until it eventually becomes zero and a deterministic model has been found. However, I have shown elsewhere[118] that this (maximum-state or best-approximation) admissible model, with a truly random source, will have on average about the same number of states as the number of observations (length of behaviour, b) and is a structurally meaningless *memory* of the observations.

It is also now of interest to look at the other extreme, not the maximum-state admissible model (perfect fit) but the minimal-state admissible model with, in fact only one state. A 1-state model can predict only a constant distribution over the descriptors, say $\mu(d)$ for $d \in D$, where:

$$\sum_{d \in D} \mu(d) = 1 \tag{3}$$

and, if there are $k(d)$ d's in the behaviour b, the total loss will be:

$$P = -\sum_{d \in D} k(d) \log(\mu(d)) \tag{4}$$

which is well known to be minimized when:[126]

$$\mu(d) = k(d)/k \tag{5}$$

where:

$$k = \sum_{d \in D} k(d) \tag{6}$$

The mean expected loss under these conditions is:

$$\max_{(P/k)} = -\sum_{d \in D} \mu(d) \log(\mu(d)) \tag{7}$$

which is the (zero-order) Shannon entropy for the distribution.

Result 2. If we plot the approximation against the complexity for the admissible models we get a monotonically falling graph that intersects the abscissa (minimum loss) at about the length of the observed behaviour if it is a Bernoulli sequence, and intersects the ordinate (minimum states) at an estimate of the entropy of the observed behaviour if it is a Bernoulli sequence.

It is this first condition, of maximal complexity, that gives an operational definition of the concept of *randomness* of even a single sequence[68] within the framework of computational complexity.[67–72] It is the second measurement, of entropy, that is conventionally taken as a measure of the randomness of a sequence. Note, however, that the sequence need not be Bernoulli (zero-memory) for the first criterion (maximum complexity) to apply. The later discussion and results will show then that the shape of the entire plot of approximation against complexity for the admissible subspace, rather than just the intercepts on the axes, may be used to analyse the randomness *and the structure* of the sequence of behaviour.

In the next main section I shall first generalize the first part of the result above to show that the *expected complexity* of a randomly generated behaviour is nearly equal to the *size* of that behaviour under very weak assumptions, and, secondly, in relation to the second part of the result, I shall derive measures of approximation with certain convergence properties that give rise to an expected loss that is the *entropy* of a behaviour, again in the general case under very weak assumptions.

3 COMPLEXITY, APPROXIMATION AND PARTITIONING OF D-SETS

In considering the identification of sequential systems the observations are essentially ordered in time and it is natural to consider them to be some subset of a free monoid of atomic descriptions (as was done in sections 2.2.1 and 2.2.2). I shall make this assumption in sections 3.3.2 and 4 to focus upon some specific inference problems. However, important results may be derived before any specific

Downloaded by [University of Calgary] at 10:52 14 February 2012

structure upon observations is assumed, and, in the present section, no structure will be postulated. The concepts developed and results obtained are thus applicable to problems of identifying systems other than automata, e.g. to problems of optical pattern recognition where there is a spatial rather than a temporal coherence between observations, or to the problem of reconciling *multiple observers* of a system where these may be only a partial order on observations from different sources.

We will take the *behaviour* of a system to be a mapping, b: $E \rightarrow D$, from a (finite) set of *events*, E, to a set of *descriptors*, D. The event space will normally have some algebraic structure, such as an order relation, upon it. Note, however, that knowledge of this structure (if it exists) is not necessary to the results of the following section (it will not be introduced until section 3.3), and that the term *event* is not intended to have necessary temporal connotations, e.g. an event might be a configuration of surface elements making up a picture and the mapping from events to descriptors might specify the reflectance and hue of each element. It is convenient to adopt Goguen's[127] neat terminology for such mappings and call the behaviour a **D-set** with E as its *support* and D as its *truth-set*. This establishes an important link to our other studies of the logic of automata,[128] and the possible logical, algebraic, topological and arithmetic foundations of automata theory. In particular it establishes a link to the wide range of results on fuzzy and probabilistic systems and the relationships between them.[117,127,129–134] Such structural considerations will be touched on only briefly in this paper (section 5.1), but form an important direction in which to extend the results.

3.1 *General Results in Complexity*

The first part of Result 2 of the previous section may now be derived with less assumptions so that it applies to the more general formulation of system behaviour given in section 2.2 and above. It is possible to generalize the enumerative technique used in Ref. 118, as Pearl[101–103] has done independently in conjunction with Shannon's rate distortion theory, to show that nearly all behaviours are complex given certain very weak and intuitive assumptions about the number of models with a given complexity and the number of behaviours of a given size.

Assume that the complexity of a model is an integer in the range 1 to infinity and that the number of distinct models with complexity, C, or less is $M(C)$. Then we may show that the mean complexity of a complete set of models up to and including those of complexity C_{max} is itself of order C_{max}.

Result 3. If $M(C)$ grows at least exponentially with C then the ratio of the mean value of C, C_{mean}, to C_{max} in a complete set of models of complexity up to and including C_{max} is asymptotic to 1.

Proof. The mean complexity is given by

$$C_{mean} = \frac{1}{M(C_{max})} \sum_{C=1}^{C_{max}} C(M(C) - M(C-1))$$

$$= C_{max} - \sum_{C=1}^{C_{max}-1} M(C)/M(C_{max}) \qquad (8)$$

Now if $M(C)$ grows at least exponentially with C then we have:

$$M(C)/M(C-1) \geqq A \qquad (9)$$

for some constant A, so that:

$$\sum_{C-1}^{C_{max}-1} M(C)/M(C_{max}) \leqq 1/(A-1) \qquad (10)$$

So that:

$$C_{mean} \geqq C_{max} - 1/(A-1) \qquad (11)$$

Hence, the ratio of C_{mean} to C_{max} is asymptotic to 1.

Exponential growth is a common feature of most model sets, such as automata, since the addition of one more state multiplies the number of possible models by at least a constant. However, it is also possible to relate the need for exponential growth in the set of models with a similar rate of growth in the set of possible behaviours. Take the *size* of a behaviour to be the number of atomic descriptors necessary to describe it, i.e. the number of events in the behaviour, and let the number of distinct behaviours of size S be $B(S)$. Suppose now that for deterministic modelling it is impossible for a given model to be an exact fit (zero approximation) to more than one behaviour, i.e. there is a mapping from models to behaviours, not necessarily $1-1$. Thus we must have that the maximum complexity of models necessary for behaviours of size S is bounded by:

$$M(C_{max}) \geqq B(S) \qquad (12)$$

If we now suppose that $M(C)$ and $B(S)$ are both similar functions of C and S respectively such that Eq. (12) implies:

$$C_{\max} \geqq S - k \tag{13}$$

where k is a constant. Then, if all behaviours are equally likely, Result 3 implies:

$$C_{\text{mean}} \geqq S - k' \tag{14}$$

where k' is a constant, provided $B(S)$ grows at least exponentially with S.

The *size*, defined in this way, of a **D-set** of N_e events over N_d descriptors is $(N_d)^{N_e}$ showing the required exponential growth. As noted previously, the action of a new element to a model will usually increase the number of models by at least N_d also. For many cases the rate of growth of models with complexity is a polynomial in N_d and C, times an exponential term of the form N_d^C so that Eqs. (13) and (14) do apply. Thus the result expressed in (14) is of wide applicability, loosely expressed:

Result 4. The mean complexity of model required over a uniform distribution of behaviours of a given size is asymptotically proportional to the size of behaviour provided the number of distinct models and the number of distinct behaviours grow in a similar fashion with respect to complexity and size, respectively, at least exponentially.

3.2 *Approximation between* **D-sets**

Having taken a behaviour, b, to be a **D-set** represented as a mapping from events to descriptors, $b: E \rightarrow D$, we might now assume that a modeller of the observed behaviour also produces some behavioural **D-set**, $m: E \rightarrow D$, as an attempt to represent b. A deterministic modeller would produce a single, unique **D-set** and one could ask whether it was identical to b. However, if the class of models available was such that identity was not possible then it would be necessary to have some measure of the extent to which m *approximates* b. One obvious measure is the total number of events on which m and b disagree:

$$N(b, m) = \sum_{e \in E} (1 - \lambda(b(e), m(e))) \tag{15}$$

where λ is a two-argument function that takes the value 1 if its arguments are equal and 0 otherwise.

$N(b, m)$ is actually a *distance measure*, i.e. we can show:

$$N(x, x) = 0 \tag{16}$$

$$N(x, y) = 0 \Rightarrow x = y \tag{17}$$

$$N(x, y) = N(y, x) \tag{18}$$

$$N(x, y) + N(y, z) \geqq N(x, z) \tag{19}$$

So that it is possible to speak of the measure of approximation as being the "distance" of the model from the behaviour. The result is dependent only on λ itself being a distance measure and hence generalizes to weighting schemes other than the simple one given above. If the **D-set** also had the structure of a monoid then the measure N could be seen as closely related to measures of *string approximation*[135] used in studies of text editing[136] and the determination of genetic ancestors.[137,138]

Measures of approximation such as N would be appropriate to a modeller that proposed just one behaviour to test against the observed behaviour. For example, in the context of modelling probabilistic automata, the modeller might put forward the behaviour having *maximum likelihood*. However, in general, an acausal modeller would propose not just one particular behaviour but rather a *set of possible behaviours*, and we need a measure of approximation that gives a distance from a set of behaviours to the observed behaviour. The minimum distance of one object from a set is one well-known extension, but fails in this case because a modeller could generate all possible behaviours and hence ensure zero distance. If, a distribution over the set of proposed behaviours is also given, however, then the mean distance of the modeller's proposed behavioural **D-set** from that actually observed would seem to be a suitable measure.

There is an alternative, but equivalent, viewpoint that throws new light on the problem. A distribution over possible behaviours is equivalent to a set of distributions over descriptors, one for each event. The modeller can then be seen to be proposing for each event not a predicted descriptor but instead a *distribution* over possible descriptors. This move from (maximum likelihood) deterministic predictions to so-called subjective probabilities has been studied both theoretically and experimentally in recent years[91–100] in order to elicit more information from human beings and to provide formal foundations for *subjective probability theory*.[139–145] It is possible to use the techniques and results developed in these studies directly in

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

relation to the current problem of acausal system identification. Indeed, the developments reported here might be seen as an extension of subjective probability theory to *sequential processes*.

Thus suppose now that the modeller proposes, not a **D-set** of descriptors, but rather a set of distributions over descriptors, a mapping, $\mu: E \rightarrow [0,1]^D$ from events to a product space of numbers between 0 and 1 that sum to unity. I shall write $\mu(e,d)$ for the proposed value assigned to descriptor d at event e—we have:

$$\sum_{d \in D} \mu(e,d) = 1 \tag{20}$$

It is simple to extend the measure of approximation N (Eq. 15) to apply to these distributions by averaging the value as previously defined over the distributions. Let:

$$NE(b,\mu) = \sum_{e \in E} \sum_{d \in D} (1 - \mu(e,d)) \lambda(b(e),d) \tag{21}$$

If the $\mu(e,d)$ were in fact used as generating probabilities to generate a single **D-set** at random to match against the behaviour, then NE would be the expected number of errors.

The table below illustrates the modelling process now envisaged and the calculation of the measure NE. For comparison a maximum likelihood proposed behaviour is also given and N is calculated.

Event	:	1	2	3	4	5	6	7	8	9
Behaviour	:	*A*	*B*	*C*	*A*	*A*	*B*	*B*	*C*	*C*
Model Distributions	*A*:	0.1	0.1	0.2	1	0.1	0.4	0.4	1	0.5
	B:	0.2	0.4	0.1	0	0.2	0.6	0.6	0	0.5
	C:	0.7	0.5	0.7	0	0.7	0	0	0	0
Max. Likelihood		*C*	*C*	*C*	*A*	*C*	*B*	*B*	*AA/B*	

so that

$$NE = 0.9 + 0.6 + 0.3 + 0 + 0.9 + 0.4 + 0.4 + 1 + 1 = 5.5$$

and

$$N = 1 + 1 + 0 + 0 + 1 + 0 + 0 + 1 + 1 = 5$$

Thus the modeller proposes for event 1, not the maximum likelihood prediction C, but instead the distribution (0.1, 0.2, 0.7) over (A, B, C), *i.e.* $\mu(1, A) = 0.1, \mu(1, B) = 0.2$, etc.

This formulation is interesting because it closely resembles the procedures used by de Finetti[58] and Savage[96] to elicit subjective probabilities from human subjects. De Finetti notes that if a target sequence is generated by a Bernoulli source and the subject gives a vector of numbers representing a distribution over possible symbols at each occurrence, then there is a loss function, that, when minimized by the subject, leads to him giving true probabilities. This is in our present terminology:

$$SE(b,\mu) = \sum_{e \in E} \sum_{d \in D} (\lambda(b(e),d) - \mu(e,d))^2 \tag{22}$$

i.e. the sum of the squares of the differences between the proposed distributions and the actual event "distribution" (1 for the event which occurred and 0 for each of the others). Savage shows the same property for an alternative loss function:

$$LE(b,\mu) = \sum_{e \in E} \sum_{d \in D} -\lambda(b(e),d) \log_2(\mu(e,d)) \tag{23}$$

i.e. the sum of minus the logarithms of the components in the distribution of the elements that actually occur in the behavioural **D-set**.

The convergence properties of these loss functions is readily demonstrated by assuming that the descriptor at event e is itself generated probabilistically by the same generating probabilities $p(e,d)$ and proving that the minimum expected loss occurs where $\mu(e,d) = p(e,d)$. It has been shown[93, 94] that there is an infinite family of such loss functions with the convergence property that a subject minimizing them is forced to give true probabilities in a probabilistic situation. De Finetti[58] showed this happened experimentally and the procedure has been used to assess "good probability assessors" in meteorology[95, 97] and to get maximum information about students' knowledge in multi-choice examinations.[99] Pearl[100] has recently given more meaning to the various measures that may be used to elicit subjective probabilities by relating them to possible hypotheses that the subject might make about the distribution of future payoffs in what, to him, is a gambling situation. For example SE corresponds to an exponential fall in future expected payoffs and LE corresponds to the slower decay of a Cauchy density. The original measure proposed, NE of Eq. (21) does *not* lead to the optimal modeller giving true probabilities, but is instead minimized by the modeller who gives *maximum-likelihood* estimates in a probabilistic situation, i.e. a distribution having the value 1 for the most likely event

and 0 for all the others. Hence, it again corresponds to well-defined and well-known pattern of decision-making behaviour.

The most striking difference between *SE* and *LE* may be seen by contrasting them on the example given previously where $E=5.5$:

$$\begin{aligned} SE &= (0.9^2+0.2^2+0.7^2)+(0.1^2+0.6^2+0.5^2) \\ &\quad +(0.2^2+0.1^2+0.3^2)+(0^2+0^2+0^2) \\ &\quad +(0.9^2+0.2^2+0.7^2)+(0.4^2+0.4^2+0^2) \\ &\quad +(1^2+0^2+1^2)+(0.5^2+0.5^2+1^2) \\ &= 1.34+0.62+0.14+0+1.34+0.32+2 \\ &\quad +1.5 \\ &= 7.26 \end{aligned}$$

$$\begin{aligned} LE &= -\log_2(0.1)-\log_2(0.4)-\log_2(0.7) \\ &\quad -\log_2(1)-\log_2(0.1)-\log_2(0.6) \\ &\quad -\log_2(0.6)-\log_2(0)-\log_2(0) \\ &= 3.32+1.32+0.51+0+3.32+0.74 \\ &\quad +0.74+\infty+\infty=\infty \end{aligned}$$

The logarithmic measure will not tolerate the situation where an event is given a valuation of zero but then occurs—the error then becomes infinite, whereas both *NE* and *SE* give large but finite errors in this situation. The logarithmic measure is also distinguished in that it depends only on the valuation given to the event which actually occurred regardless of the distribution over the other components. This has been taken by some writers as a desirable feature although the argument seems dubious and there are more meaningful considerations that make the logarithmic measure attractive (see following section).

One important aspect of the move earlier in this section from the concept of a modeller proposing a distribution over possible behavioural **D-sets** to the concept of its proposing a set of distributions is that the loss measures may be regarded as having a component associated with each event. The overall loss is, in all three cases, the sum of the losses associated with each event. The component added for event may be described as the *surprise* caused by that event. All three measures agree that the surprise caused by an event given the valuation 1 which actually occurs is zero (e.g. event 4 in the example). They give varying weights to events which would occasion little surprise (e.g. event 3) or much surprise (e.g. event 1) and, as noted, the logarithmic rule expresses infinite surprise at an event that occurs when the valuation given to it is zero. This valuation of "surprise" is consistent with the model of decision-making based on "*potential surprise*" proposed by the economist Shackle,[146,147] and is particularly useful in on-line learning algorithms where a marked increase in the rate of surprise may be used to indicate the need for the recomputation of the model.

3.2.1 *Entropy—the expected loss for probabilistic behaviour* One can avoid the premature use of evocative terms such as *subjective probabilities* for the distributions proposed by the modeller in the previous section, preserving methodological neutrality until results, under certain circumstances, prove the terms justified. The theoretical and experimental studies of de Finetti, Savage *et al.*, indicate that *if* the actual event sequence is probabilistically generated then a modeller that is optimal (in the sense of minimizing the poorness-of-fit measures, *SE* or *LE*) will be forced to propose the actual generating probabilities of events. This result is an important link between *subjective* and *physical* or *frequentist* accounts of probability theory. It is equally important as a link between our general approach to system identification and probabilistic modelling. However, the measures defined in the previous section and the identification techniques based on them do *not* in themselves entail a hypothesis of probabilistic acausality. The fact that they behave meaningfully and well when used with probabilistic systems is clearly desirable, even essential, but there is no converse argument that they are based on a hypothesis of probabilistic behaviour in the system modelled.

Clearly, we may now expect to obtain results for probabilistic modelling (optimality of identification techniques, decision criteria for selecting amongst admissible models, etc.) which do not necessarily apply in more general cases—indeed are not meaningful unless further hypotheses are made about the more general case. Clearly also, there are few hypotheses comparable in power and significance to that of a probabilistic generator. In the experimental studies, I have taken examples of asynchronous systems modelling where no probabilities are definable but it is possible to obtain weaker, structural rather than numeric, results for identification techniques based on the measures of approximation defined. An example of nonprobabilistic acausality will be given in section 4.3 where several samples of the behaviour of a deterministic system are identified—the acausality arising through the sampling process and having no numeric, probabilistic significance in the model. Such examples are important in demonstrating the generality of the approximation measures outlined and also

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

illustrate the role of the probability *logic*[148,117] underlying probabilistic models in representing more general acausal phenomena: Kalman[149] has shown that whilst linear operational techniques may be applied to sampled-data or discrete systems separately, the most appropriate representation of a sampled-data, discrete system is pseudo-probabilistic; I have shown elsewhere[128,150] the wider roles of probability logics in modelling *possibility* and *eventuality*; and close links have been established[151-153] between probability logics and *modal logics*[154-155] of possibility, necessity and time.[156]

However, it is also of interest to determine what happens in a truly probabilistic situation when a modeller does manage to propose precisely the optimal distributions that he is forced to converge towards when minimizing the loss functions *SE* and *LE*. Suppose the actual probability of occurrence of descriptor d at event e is $p(e,d)$ and the distributions proposed by the modeller are such that $\mu(e,d) = p(e,d)$. The expected values of the loss functions of equations 21 through 23, *NE*, *SE* and *LE*, respectively, are:

$$\widehat{NE}(b\,p) = \sum_{e\in E}\sum_{d\in D} p(e,d)(1-p(e,d)) \qquad (24)$$

$$\begin{aligned}\widehat{SE}(b,p) &= \sum_{e\in E}\sum_{d\in D} p(e,d)(1-p(e,d))^2 \\ &\quad + (1-p(e,d))(p(e,d))^2 \\ &= \sum_{e\in E}\sum_{d\in D} p(e,d)(1-p(e,d)) \qquad (25)\end{aligned}$$

$$\widehat{LE}(b,p) = \sum_{e\in E}\sum_{d\in D} p(e,d)\log_2(p(e,d)) \qquad (26)$$

Equations 24 through 26 show that the expected value of the loss, or approximation measure, when the modeller matches the actual generating probabilities is actually an *entropy*[157] function for all three measures. In particular *LE* of Eq. (26) is the familiar Shannon entropy whose significance in inductive inference has been emphasized by Watanabe.[158] The coincidence in values of *NE* and *SE* is interesting but spurious since, as noted previously, whilst the condition $\mu(e,d) = p(e,d)$ gives a minimum for *SE* and *LE*, the corresponding condition minimizing *NE* is:

$$\mu(e,d) = \begin{cases} 1 \text{ if } d = d' \\ 0 \text{ otherwise} \end{cases} \qquad (27)$$

where d' is some descriptor such that:

$$p(e,d') = \max_{d}(p(e,d)) \qquad (28)$$

a maximum likelihood estimate. Thus De Finetti's quadratic loss function *SE* of Eq. (22) has the same expected loss as the more obvious definition of *NE* in Eq. (21) (expected number of errors) but forces actual probability estimation on the modeller which *NE* does not. It is also interesting to note that *NE* and *SE* are the expected value of $cont(e\&d)$ and *LE* is the expected value of $inf(e\&d)$, linking these results to theories of semantic information and logical probability.[8-13,60-63]

To give physical significance to the values $p(e,d)$ within the current framework (in which the events in E have so far been regarded as unique) we need to hypothesize some structure on the event space such that many different events may be regarded as equivalent. Then the descriptors at those events may be regarded as a sample space for some process generating descriptors with a probability distribution $\mu(e,d)$ (where $\mu(e,d) = \mu(e',d)$ for equivalent e and e'). Such an equivalence relation between events makes it possible for events, under the equivalence relation, to be recurrent, and hence for a modeller to determine the generating probabilities. If there is more than one equivalence relation (i.e. all events are not equivalent to one another) then there is also some *structure* present to be determined by the modeller. The following section discusses this problem of the determination of this structure.

3.3 *Structures on* **D-sets**

We have come remarkably far without assuming any structure on the **D-sets** representing behaviour and proposed by modellers. This has been quite deliberate since I wished to effect a clear separation between: (*a*) concepts of complexity and randomness (section 3.1); (*b*) measures of approximation between behaviour and model (section 3.2); and (*c*) the modelling of sequential machines (this section). It will be noted that the developments in sections 3.1 and 3.2 are independent of one another and neither requires such postulates as: the behaviour is that of a sequential machine; events are ordered in time; etc. The concept that a *random* behaviour is one whose *complexity* is nearly equal to its *size*, and the result that the expected loss in approximating a random event is an *entropy* for it, have both been obtained under very weak assumptions and constraints. In particular neither involves notions of

causality or of *computational complexity* in terms of sequential machines. In the context of the general formulation of the problem of identification given in section 2.2, the first result relates to the intercept of the plot of approximation against complexity for admissible models with the axis of zero approximation, and the second result relates to its intercept with the axis of zero complexity and to the significance of the approximation at different complexities. Thus the two results are linked through the concept of an admissible subspace of models.

I emphasize the general status of these results because they will clearly have a role to play in any more specific formulations of the identification problem. In a sense they will be *re-discovered* for each class of behavioural and model structures, and a major virtue of a general systems theoretic formulation is to guide this re-discovery and allow the fundamental status of these concepts and results to be seen independently of the class of problems for which they are obtained. Conversely, the essentially *structural* role of notions such as causality can be seen more clearly if they are separated from the numerical notions of complexity and entropy and only introduced at this late stage.

3.3.1 *The "Rules of the Game"* What then are the rules of any further constraints upon the behavioural **D-set**, that it is *ordered*, *monoidal*, etc.? In fact, they are best viewed as *rules of the game* whereby the modeller is allowed to use the descriptors attached to some subset of events as a basis for proposing distributions that apply to others. Consider the two extremes: the modeller has to propose distributions completely *a priori*; with complete knowledge of events the modeller can propose precisely the actual behaviour. Both these extremes are trivial and in any realistic problem formulation the modeller will have *partial knowledge* about some events that will influence his proposals about others.

I call the constraints upon the use of partial knowledge *rules of the game* because the modeller can usefully be viewed as playing a game against an opponent, or at least being constrained by a referee—he has to state what algorithm, within the rules of the game, he is using to propose a distribution and then his opponent or the referee evaluates it and reports back the loss, or degree of approximation, achieved. This concept is useful because in practice, since we are concerned with the problem of formulating a model for an already known behaviour, of identifying its underlying generative structure, the modeller actually knows all the behaviour and it is what he is allowed to *use* that is constrained. For example, if events are ordered (e.g. in time) and the modeller in proposing a distribution over descriptors for a particular event is allowed to use only the value of descriptors for events properly less than that one, then it is said to be *causal*.[159–161] If there is a metric topology on the event space and the modeller in proposing a distribution for a particular event is allowed to use only the value of descriptors for events in a neighbourhood of specified size about that event then it is said to be *local*. If the algorithm used by the modeller in determining its proposals for each event is independent of the particular event then it is said to be *uniform*.

Non-uniform modelling can be justified if the event space has known inhomogeneities, e.g. boundaries in optical pattern recognition. *Uniform local* algorithms are appropriate to such operations as contour extraction in scene analysis[162, 163] and has a natural interpretation in terms of Lie groups.[164, 165] They have been studied, for example, by Minsky and Papert [166] in determining the power of "diameter-limited perceptrons". *Uniform local causal* algorithms correspond to so-called *finite-memory*[167, 168] algorithms, and have been studied as *definite events*,[169] locally-testable languages,[170, 171] etc.

Local and finite-memory algorithms must not be confused with *finite-state* algorithms. The notion of state, as developed by Birkhoff,[172] Markov,[173, 174] Zadeh,[175, 176] Salovaara,[177] and others, is a system concept of major importance in its own right and quite distinct from that of "finite memory". There is an ontological distinction between them in that the "finite memory" refers to a bounded locality accessible in the environment, in the observed behaviour, whilst the "finite-state" refers to a limited storage capacity in the algorithm, or model.

States are best introduced into the current discussion by considering them as (conceptual) *tokens* that may be placed on events as an aid to proposing distributions. A *finite-state* algorithm has only a specified finite number of distinctions between tokens that it may make. The placing of tokens is governed by rules of the game similar in their possibilities to those for proposing distributions. However, the placing of tokens and the proposing of distributions is now allowed to depend not just on the description at a specified set of events but also the tokens placed at them.

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

The terms *local* and *causal* may clearly be applied to the algorithm for placing tokens at events as before. However, an interesting new distinction arises since the descriptor actually at the event may, or may not, be allowed to influence the token placed there (it would clearly be trivial to allow it to influence the proposed distribution for descriptors there). If the algorithm for allocating a token to an event is dependent only on descriptors and tokens at other events and gives a definite result it is said to be *deterministic*. If it gives a definite result only when the descriptor at the event is also specified it is said to be nondeterministic but *observable*; otherwise it is just *nondeterministic*. Thus, in the context of causal systems, a nondeterministic but observable algorithm cannot predict the next state precisely but knows what it is when the next descriptor is known, i.e. it can always specify the current state precisely, not just as a distribution.

The concept of state is normally associated with causal processes, i.e. time sequences. If we specify that the events are either chain-ordered or incomparable so that the behaviour may be regarded as a submonoid of the free monoid of descriptors, and that a local neighbourhood of an event are the events immediately next to it in its chain. Then a *local causal deterministic* algorithm is a deterministic sequential machine or *automaton*, and a *local causal observable* algorithm is an observable probabilistic automaton.

The role of states in general may be seen as that of *extending the locality* open to a local algorithm. Tokens may be placed to indicate features of the event space that are not apparent from inspection of the allowed local neighbourhood. For example, in pattern recognition the *global property* of being part of an oriented line, as opposed to the local property of being an oriented line segment, requires the placing of tokens to compute. It is the defect of the perceptron that it cannot place such state tokens that Minsky and Papert[166] studied, demonstrating the resultant weakness with respect to the determination of global properties.

It will be noted that, whereas in a causal system there is an unambiguous sequence for the placing of tokens, the ordering in time, this does not exist in general and problems may arise in, for example, the spatial assignment of tokens such that many, or no, final assignments are possible. It is interesting to speculate that this may correspond to some of the phenomena of spontaneous change in perception[178] whereby, for example, a figure will be first perceived as one object and then as another—the assignment of local features to global percepts being ambiguous and a number of different, but self-consistent, assignments being possible.

Informally also the concept of states as conceptual tokens placed in the environment is attractive because it is apparent that many problems of sequential inference are avoided in human and animal societies by placing actual *physical tokens* in the environment. The internal states necessary to an algorithm solving an essentially sequential inferencing problem may be seen as compensating for the lack of such physical signposting in the actual environment. Since it will be clear that even simple sequential inferencing tasks can require massive computation it may well be that the mathematical models of induction thus generated are *not* a reasonable representation of human thought processes. *We mark the environment* to make inferencing simple and to avoid the need for internal states (short-term memory) as much as possible.

3.3.2 *Probabilistic automaton structures* If we now focus down upon local causal algorithms over descriptor monoids, i.e. probabilistic automaton, it would be useful to derive some overall convergence results similar to those for deterministic automaton modelling using the Nerode equivalence. These are[118] primarily:

1) The number of states in the model is a monotonic non-decreasing function of the length of the observed sequence of behaviour. Thus the model cannot become simpler with further observations;

2) The number of states in the model of a sequence of behaviour generated by an M-state deterministic automaton cannot exceed M. Thus the model cannot grow more complex than the generating system.

The situation is far more difficult to analyse for probabilistic modelling because there is now no well-defined *best model* but instead an admissible subspace of possible models. However, some comparable results may be obtained. Any probabilistic automaton may be regarded as outputting a distribution over possible outputs at each step (this is not identical in concept to an ensemble of all possible transitions, but rather more artificial since we are following the state-distributions of a single automaton). Consider the problem of how a modeller might actually come to match a sequence of distributions. The two sources of difficulty are:

a) Since the sequence is not Bernoulli and the distributions change from event to event the

modeller must be able to locate himself in the sequence, i.e. the events with different distributions must be *observable*. This is similar to the situation in deterministic modelling where no modeller can discriminate between two different structures if their reduced, observable forms are isomorphic. In practice, for a probabilistic model, this condition implies simply that, even though from a state we can predict only the probability of the next state, after the transition the output must be sufficient to indicate the actual state. Thus in analysing the match between source and model we need only consider the reduced form of the source—this appears in result (2) above as the number of states in a model cannot *exceed* those in the source rather than becomes eventually equal to their number.

b) The distributions themselves are not the actual outputs and an indefinitely large sample of actual outputs is necessary in order to estimate them. Thus the distributions in any *transient* behaviour of the source automaton cannot be precisely estimated, only those in its recurrent behaviour.

Combining these two factors we can see that it is realistic to consider matching precisely in a model only the *recurrent* behaviour of the *reduced* form of a probabilistic automaton. The recurrency is an additional constraint compared with deterministic modelling and clearly a reasonable one in the circumstances. Conceptually the modeller of an observable sequence of distributions is applying a Bernoulli sequence modelling strategy to each distinct distribution. However, he has *both* to discover the observation algorithm *and* estimate the distributions.

Consider the formulation of the identification problem given in section 2.2.2 with the identification space, $(D^*, M, \leqq, f)$, where: M is the space of probabilistic Moore automata in reduced form; $m \leqq n$ if m has less states than n; the value $f(b) = \leqq_b$ being defined for a sequence of behaviour b by $m \leqq_b n$ if $LE(b, \mu_m) \leqq LE(b, \mu_n) + \varepsilon s$ where μ_m and μ_n are the distributions arising from modelling m and n respectively, s is the length of b and ε is a small tolerance to allow for statistical fluctuations (an alternative f may be similarly defined based on SE instead of LE). Now consider the behaviour of the admissible subspace for a sequence of observations of increasing length of a finite state probabilistic source.

A result equivalent to (1) above could be that the maximum number of states in the admissible models is monotonic non-decreasing. However, this is not so because in the short term the particular sequence generated may be such as to justify complex models. However, there is a result equivalent to (2):

2′) For any ε, for a given probabilistic source, the maximum number of states of an admissible model cannot eventually exceed that of the source as the sequence of observation increases. This follows because the properties of LE are such that averaged over a long sequence of recurrent states any modeller cannot do better than put forward precisely the distributions of the reduced form of the source and hence the source itself will have at least as low a value of LE as any model with a greater number or states.

We still cannot show that the maximum number of states is precisely that of the source, even in reduced form, because the modelling procedure cannot necessarily identify the transient behaviour of the source. Our definition of $\leqq_b$ based on LE and s means that in the long term the transient behaviour will have a decreasingly small effect so that eventually the admissible models neglect it. The maximum number of states in an admissible model will then correspond to the number of states in the recurrent part of the automaton generating the observed behaviour (there is clearly no well-defined *recurrent part* in general since we may enter different recurrent parts after the same initial transient). What of the admissible models with less than the maximal number of states? These correspond to the "lumping" of states in a Markov process and will inevitably give higher values for the entropy of the process and hence LE. They are best approximations to the source in the sense that they minimize the deviation in behaviour from that of the actual source.

In the next section some experimental results on an actual computer implementation of a probabilistic automaton modelling algorithm are given that illustrate the points made in this section.

4 SOME COMPUTATIONAL STUDIES

The approach to the problem of identification developed in the previous sections is intended to be completely operational and hence open to computer implementation and experimental study and application. Even the most general formulation of section 2.2 is mathematically precise in terms of order relations and open to computational study for any

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

identification technique. Probably the most interesting level at which to undertake such computational study at the present time, however, is that of probabilistic automaton inference. This is currently an unsolved problem with many interesting recent case studies, generally in terms of probabilistic grammatical inference,[73–76] that is also of great practical importance for the analysis of real-world data such as human and animal behaviour patterns.[179–182]

The main reason that the problem of inferring the structure of a discrete probabilistic system from its behaviour has not been solved (in the sense that the equivalent problem for deterministic systems has) is that discussed in section 2, that the nature of a "solution" is not well defined—we have had no decision procedure to determine that a particular probabilistic automaton is a "correct" representation of a given sequence of behaviour. Thus, the most successful studies of the problem have tended to be those that have concentrated on the methodology of *approximation*, of deciding when a particular automaton or grammar is *reasonably correct*. In particular the recent studies of Maryanski,[76] part of a programme of research directed by Booth,[74,75] provides a series of challenging problems and results for other workers. The concepts of an admissible subspace (section 2.2), and the specific measures of approximation (section 3.2), developed in this paper provide the basis for a formulation and solution of the problem of probabilistic structure inference that is precisely defined, and hence *complete* in a sense that previous formulations have not been, and it is clearly of interest to evaluate the operation of the techniques on actual data.

The following sub-section describes ATOM, a suite of programs embodying the identification techniques described in this paper. A reanalysis of some of Maryanski's data using ATOM, is given in Ref. 90, and a variety of informal examples in Ref. 89. The examples in the paper have been selected both to demonstrate the basic techniques and also to illustrate the hierarchical analysis of epistemological levels in system identification proposed by Klir[104,240] and discussed in section 5.1 of this paper.

4.1 *The ATOM Identification System*

The identification system to be described is one in which the class of models M is that of finite-state probabilistic automata with the relationships, $\leqq$ and $\leqq_b$, defined as in section 3.3 by the number of states in the model and the logarithmic approximation measure, LE, respectively. The computational algorithms to determine the admissible models are one of a suite of programs called "ATOM" written in the interpretive, string-handling language BASYS[183,184] on a time-shared PDP11/45 (the algorithms have also been used in FORTRAN on a PDP10 with KA10 processor and run some 50 times faster). ATOM provides facilities for interactively entering observed data and forming on-line predictions from models, and so on. However, for the automata modelling studies it is generally used in restartable background batch mode since computational runs of hundreds of hours may be required.

A behaviour to be modelled is input to ATOM as a string of arbitrary character strings separated by spaces or end-of-line terminators. Thus:

```
MARY HAD A LITTLE LAMB
ITS FLEECE WAS     ?
```

is a sequence of behaviour consisting of 9 symbols, and:

```
I=2
P=A(I)
J=P/I+7
```

is a sequence of behaviour consisting of 3 symbols. This acceptance of free format strings is particularly helpful in some examples such as natural language processing and automatic programming.

ATOM assumes that all the symbols are automaton *outputs* unless it is separately informed that a certain set of symbols are *inputs* and/or another set are *delimiters*. All the modelling schemes treat these two classes in a similar fashion: *inputs* are not brought into the string approximation measurement, i.e. one does not evaluate the extent to which input symbols are predicted correctly; *delimiters* are taken to indicate that the string before the delimiter may not be used to predict that after it—in the automaton models a delimiter causes a reset to a single state of the model. Otherwise both inputs and delimiters are treated as any other symbols in the string of behaviour. Note that the availability of delimiters enables separate samples of behaviour (separate sentences say in a grammatical inferences problem) to be freely concatenated together, separated by delimiters, to form a single sequence of behaviour. Note also that the modelling process does not necessitate inputs and delimiters being specified

Downloaded by [University of Calgary] at 10:52 14 February 2012

in this way. If they are then the computation is faster, but if they are not then their nature may be inferred from the results—i.e. inputs are "outputs" that cannot be predicted and delimiters are those which appear as a general reset—examples of such inferences will be given later.

The automation identification subprogram in ATOM generates, for a given behaviour, the admissible subspace of either Mealy or Moore probabilistic automata, as requested, commencing with 1-state models (Mealy) or k-state models Moore—where k is the number of different symbols in the behaviour). The actual output of the programme is thus the set of best-approximation 1-state models, the set of best-approximation 2-state models, etc. The search ceases when no more admissible models are found, but in practice this condition rarely arises since the search space for larger models grows rapidly with the number of states and the programme is terminated by lack of time rather than by completion of the search. However, since the simpler admissible models are output first, the modelling is always complete up to models with the number of states at which it was terminated.

The search procedure is essentially simple because *only the space of non-deterministic automata has to be searched*, not that of probabilistic automata, i.e. the transitions are initially regarded as being only present or absent. When a non-deterministic model of the behaviour has been generated the actual transition probabilities are filled in from the relative frequencies of the transitions in the particular model with the given behaviour. This is legitimate because these values are known to minimize the *LE*. The value of *LE* for the model/behaviour pair is then calculated and used to ascertain the approximation relative to previous models generated. If the approximation is the same, or better, than that of the best models previously formed then the new model is added to the set of potentially admissible models. Any models with the same number of states but a poorer fit on both criteria are discarded. The search then continues. Whenever the models with a given number of states has been searched then the remaining best models with that number of states are filed as being admissible. The normalized values of *NE* (assuming maximum-likelihood estimates) and *LE* are filed with the model.

The generation of models is basically an exhaustive enumeration of all possible observable non-deterministic automata. However, some care is necessary to avoid duplication and to take advantage of any structural features of the sample behaviour (e.g. some symbols never following other symbols). Models are generated using the actual behaviour to fill in state transitions. The initial model is a 1-state automaton and, if N-state models are being searched, N is taken as a bound on the number of states. The initial state has to be the first and only state. Each symbol in turn in the behaviour is then examined. If it corresponds to an existent transition no action need be taken. If there is no transition corresponding to it then one is entered to the first state and a marker placed on a stack. The state is then advanced to its next value and the next symbol checked.

Eventually a model has been formed and may be evaluated for *LE*. A backtrack is then made by taking off the stack the last transition entered and, if it is to state k, changing it to be to state $k+1$ and continuing as before. However if state k was a new state then it is removed and backtracking performed again, or if k was the last state and not new, and k is less than the allowed maximum number of states, then a new state is added and the transition entered to this. Eventually backtracking is no longer possible and all models with the allowed number of states have been generated without duplication and without considering transitions not necessitated by the behaviour being modelled.

"Delimiter" symbols are taken to cause a reset to a single state (usually the initial state since strings with delimiters normally commence with one). "Input" symbols are not taken into account in the calculations of *LE*. The following sections contain examples of ATOM automaton modelling in action.

4.2 *Identification of a Five-State Stochastic Automaton*

Figure 1a shows a sequence of 200 descriptors generated by a 5-state binary stochastic source (read in the natural order of English text, from left to right, line to line). This was input to ATOM and the admissible subspace of models computed up to 6 states based on a normalized *LE* measure of approximation. A deterministic model was also computed. Figure 1b shows the ATOM output of 1 through 6 state admissible models. The first line gives *NE* (as fraction and percentage); followed by PLOGP (normalized *LE*); the type, number of states and number in order of search of the model; and the name of data file. The following lines describe the state transitions of the model: for each

state, the transition under each descriptor and the number of times it occurs (in brackets), e.g.:

$$1{:}\,B\rightarrow 1(34)\;A\rightarrow 2(83)$$

means that in state 1 a B leads to state 1 34 times and an A leads to state 2 83 times. The final line after each model shows the total number of models searched—note that *no* 6-state models exist that are better then the 5-state best.

Figure 1c is a plot of approximation against complexity (PLOGP or *LE* against number of states) for the admissible models, and Figure 1d is a set of diagrams of the models themselves with descriptor associated with the transition subscripted with the number of times the transition occurs. Note how *LE* commences near its maximum value of unity for a 1-state (Bernoulli sequence) model, corresponding to the total number of A's in the sequence being close to that of B's. It drops sharply by one half for a 2-state model that shows

FIGURE. 1 Identification of a 5-state stochastic automaton.

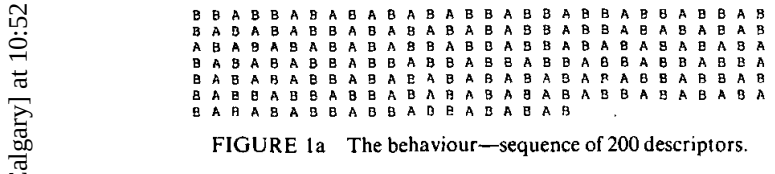

FIGURE 1a The behaviour—sequence of 200 descriptors.

```
83/200 (41.50%)     PLOGP=0.978       MEALY:1:1     FIVE.DAT
* 1: B -> 1(117) A -> 1(83)

1 MODELS WITH 1 STATES - BEST PLOGP 0.978

33/200 (16.50%)     PLOGP=0.499       MEALY:2:8     FIVE.DAT
  1: B -> 2(84)
* 2: B -> 2(33) A -> 1(83)

9 MODELS WITH 2 STATES - BEST PLOGP 0.499

33/200 (16.50%)     PLOGP=0.402       MEALY:3:51    FIVE.DAT
  1: B -> 2(84)
* 2: B -> 3(33) A -> 1(50)
  3: A -> 1(33)

85 MODELS WITH 3 STATES - BEST PLOGP 0.402

31/200 (15.50%)     PLOGP=0.385       MEALY:4:529   FIVE.DAT
  1: B -> 2(57)
* 2: B -> 3(29) A -> 4(27)
  3: B -> 3(4) A -> 1(56)
  4: B -> 3(27)

969 MODELS WITH 4 STATES - BEST PLOGP 0.385

19/200 (9.50%)     PLOGP=0.317       MEALY:5:6215 FIVE.DAT
  1: B -> 2(34)
  2: B -> 3(24) A -> 4(10)
  3: A -> 1(33)
  4: B -> 5(50)
* 5: B -> 3(9) A -> 4(40)

13378 MODELS WITH 5 STATES - BEST PLOGP 0.317


216347 MODELS WITH 6 STATES - BEST PLOGP 0.317
```

FIGURE 1b ATOM analysis of behaviour—1 through 6-state admissible models.

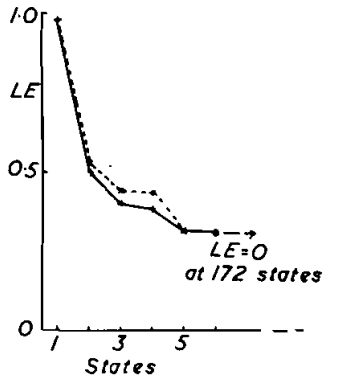

FIGURE 1c Plot of approximation against complexity (logarithmic measure, *LE*, against number of states)—dashed plot from validation sample.

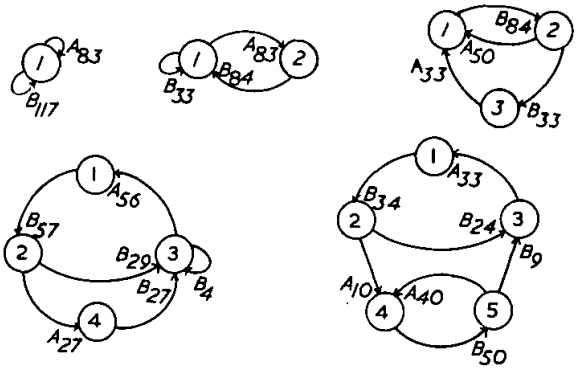

FIGURE 1d Admissible models produced by ATOM up to 5 states.

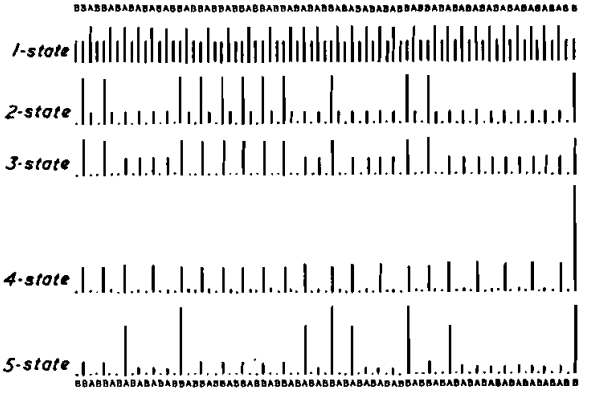

FIGURE 1e Surprise at each of first 70 descriptors for 1- to 5-state admissible models.

FIGURE 1f Independent second sample of behaviour for validation.

Downloaded by [University of Calgary] at 10:52 14 February 2012

essentially that A is always followed by B. It drops substantially again for a 3-state model that adds the constraint that BB is followed by A. Going to a 4-state model produces no significant drop but there is one in going to a 5-state model. In fact this gives an accurate model of the source which was a 2-state system generating (AB)* that had a probability of 0.2 in one of its states of transiting to a 3-state system generating (ABB)* that in its turn had a probability of 0.3 of returning to the original system from one of its states. Going on to a 6-state system produces no significant drop in *LE* and in fact *LE* reaches zero only with a (deterministic) model of 172 states.

In terms of regular events the actual sequence is of the form (BBA|BA)* and the useful models produced are:

1-state: (A*B*)* $LE = 0.978$
2-state: (BB*A)* $LE = 0.499$
3-state: (BBA|BA)* $LE = 0.402$
5-state: (BBA|BA)* $LE = 0.317$

The drop in *LE* between 3 and 5 states is caused not by an improvement in the language structure in these terms but by the decoupling inferred, that a BBA is more likely to be followed by a BBA, and a BA by BA, than the alternative possibilities, i.e. a better indication of the language structure would be ((BBA)*(BA)*)*.

Figure 1e shows the pattern of "surprise" (minus log of predicted probability for descriptor that actually occurs) for the first 70 descriptors in the sequence for each of the admissible models. It is apparent how the surprise caused by the deterministic transitions (within the two subsystems) drops to zero except for those where transitions may occur from the BBA subsequence to the BA subsequence, or vice versa. At these points the surprise at transitions within a subsystem is low whereas that at transitions between subsystems is high. Note that the "surprise" patterns shown are retrospective—they result from "replaying" the entire behaviour through the final model based on it. A different pattern would emerge if the surprise were calculated incrementally with the modeller calculating a model on the basis of the sequence so far observed. However, I shall not consider the incremental dynamics of such *learning* situations in this paper.

I have used the term *significant* for the changes in *LE* in the previous paragraph. However, apart from the visual appearance of Figure 1c where changes may clearly be seen, this term has not been justified. Some studies of mathematical significance in entropy change and uncertainty analysis has been made[185–187] but much more foundational work is needed to provide formal significance of the ATOM results, e.g. taking into account the number of models searched and the length of sequence analysed. One heuristic approach that is useful with the current system is to give ATOM a second, independent sample from the same source and get out values of *LE* for the admissible models generated from the original source, i.e. validate the models against an independent sample. This is the source of the dashed plot in Figure 1c that was obtained by evaluating the second sample shown in Figure 1f against the models of Figures 1b or 1d. Note how the 5-state, "correct", model stands out as not losing any significance whereas the dashed curve of approximation for the other models lies above that based on the sample that generated the models.

4.2 *A Transient Behaviour from a Five-State Automaton*

The test against an independent sample suggested and demonstrated in the previous section makes sense if the source is in a recurrent phase. However, ATOM also provides meaningful models of transient sequences. Consider, for example, similar behaviour to that already analysed but where the BA to BBA transition only occurs once. The sequence $(BA)^{15}(BBA)^{10}$ was input to ATOM and analysed up to 6-state models to produce the plot of

FIGURE 2 Identification of a transient behaviour.

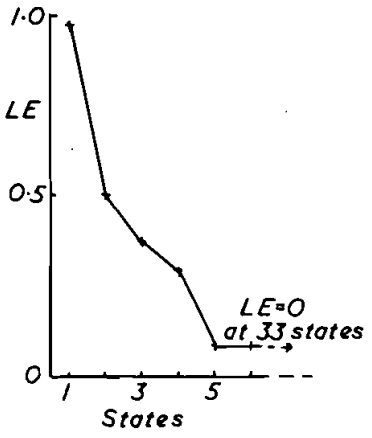

FIGURE 2a Plot of approximation against complexity (logarithmic measure, *LE*, against number of states).

Downloaded by [University of Calgary] at 10:52 14 February 2012

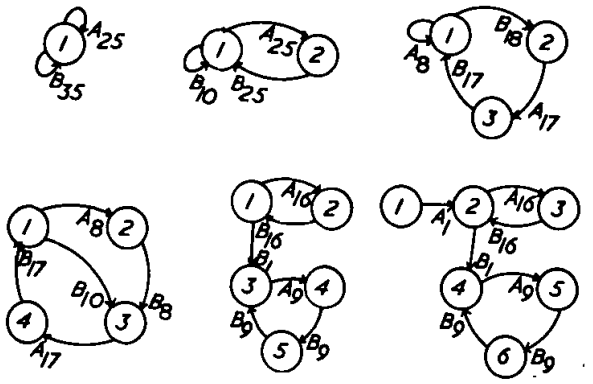

FIGURE 2b Admissible models produced by ATOM up to 6 states.

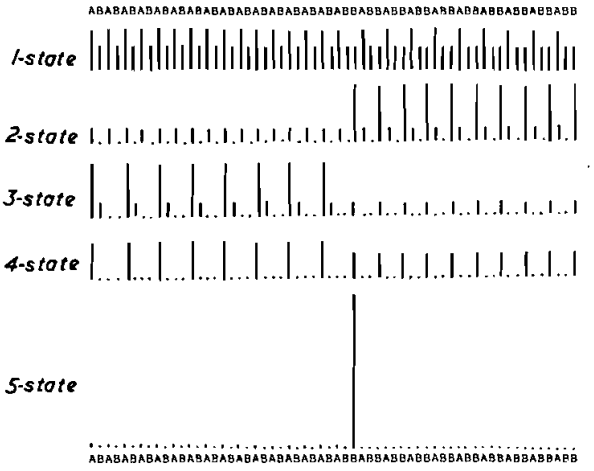

FIGURE 2c Surprise at each descriptor for 1- to 5-state admissible models.

approximation against complexity shown in Figure 2a. A turnover at 5 states is again apparent and the models of Figure 2b again show improving hypotheses as to the actual sequence structure. Note, however, in this case the final improvement in approximation at 5 states where the model actually breaks into two parts for the first time—there is only a single transition from state 1 to state 3 and there is no return from states 3, 4 or 5 to states 1 or 2. The patterns of surprise shown in Figure 2c also show up this clear change when a 5-state model is reached, with a very clear marker at the point of transition from the transient sequence to the recurrent sequence.

Thus ATOM can provide a very clear analysis of a sequence of behaviour that shows an isolated, non-recurrent change, where concepts of stationarity do not apply.

4.3 *Miscellaneous Studies—Delimiters, Inputs and Outputs*

References 89 and 90 contain a range of examples of ATOM inferencing models from behaviour in a variety of applications including: stochastic grammatical inference; the analysis of human behaviour; and the derivation of programs from their traces—this last being an application of grammatical inference studied previously[188–190] only for cases where there are no errors in the trace and use of the Nerode equivalence gives an equivalent deterministic automaton flowchart for the program trace. There are two main features of interest in these other studies that are relevant to the general formulation of this paper—the hypothesizing, or inferring, of *delimiters* and of *inputs*.

To illustrate the role of delimiters, Figure 3a shows the observed behaviour of a system as a sequence of 101 descriptors from the set {A, B, C, D}. In fact the behaviour is that of a 3-state deterministic machine generating the repetitive sequence (CBA)*, sampled in short segments of arbitrary length with the symbol D inserted as a delimiter at segment boundaries. It consists overall of 19 samples of behaviour concatenated together between delimiters.

Figure 3b is a plot of approximation against complexity for admissible models of the sequence and a turnover at 4 states is apparent. The transition diagram of Figure 3c shows the 4-state admissible model and the dominant CBA cycle is apparent. Superimposed on this are some less frequent "noise" transitions of which the most notable are those involving D since it can be seen that D may occur in any state, always leads to state 1, and from this state any descriptor may occur. The fact that D always leads to a single state shows that it acts as a *reset* input to the automaton causing a return to a single state. Since this implies in its turn that there is no memory in the automaton of the behaviour prior to D such a reset input is also called a *delimiter*. The occurrence of a D transition from any state and the exit to any state from the reset state suggests that state-independent, i.e. asynchronous, sampling is taking place. Note that the overall model is a *Moore* automaton in which transitions correspond to states: state 1 to D; state 2 to B; state 3 to A; and state 4 to C.

Figure 3d shows the surprise pattern for the initial sequence of observations from which it can be seen that D and the descriptor after it are both surprising but the rest of the descriptors, within the sampled

Downloaded by [University of Calgary] at 10:52 14 February 2012

FIGURE 3 Identification of a sampled deterministic automaton.

A D B A D A C B D C B A C B D B A C B A C D C D A C B A D
A C B D B A C B A C B A C D B A D A C B A C B A C B A D A
C B A C B A C B A D C D B A C B A D B A C D B D C B A C D
A C B A C B A C B A C D D A

FIGURE 3a The behaviour—sequence of 101 descriptors.

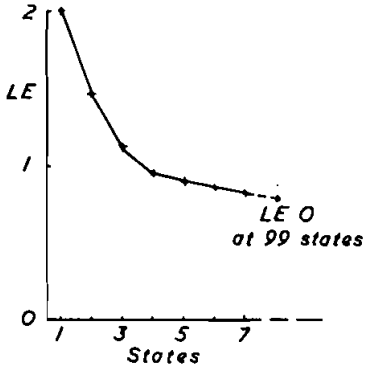

FIGURE 3b Plot of approximation against complexity (logarithmic measure, *LE*, against number of states).

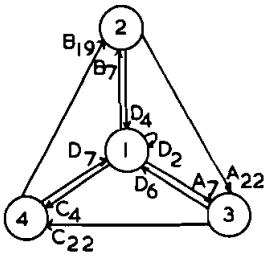

FIGURE 3c 4-state admissible model.

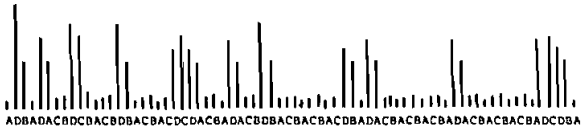

FIGURE 3d Surprise at each descriptor in initial sequence of behaviour.

behaviour pattern, are not. Thus ATOM has effectively derived the original structure of the system generating the behaviour. Note that this example involves no truly probabilistic phenomena. The sampled system is completely deterministic. The sampling itself is not random but only asynchronous and hence non-deterministic—longer samples were allowed deliberately in the later part of the sequence. As demonstrated in the previous example, the ATOM algorithm copes with non-probabilistic indeterminacy and derives the correct structure without strong probabilistic assumptions.

This example may also be used to illustrate a further feature of ATOM. Rather than inferring that a descriptor is a delimiter as a result of the ATOM analysis, we can specify in advance that it is to be treated as one. ATOM then only generates models in which that descriptor acts as a reset input, i.e. all transitions from it lead to only one state. This saves computation in searching the space of models, e.g. in deriving the 4-state model above over 10^8 4-state Mealy models were searched but if D is specified to be a delimiter this is to be reduced to 10^5 (if, additionally, only Moore models are searched this drops to 1). Hence, *a priori* structural knowledge may be traded directly for computing effort. However, if the structural hypothesis is incorrect then, as demonstrated in Ref. 118, the results obtained may be, not just an approximation, but instead totally meaningless.

The ability to specify a descriptor to be a delimiter is useful in allowing ATOM to be used to analyse behaviour that is a sub-monoid (i.e. a set of strings) rather than just a single string. If AB, CAA, and ABC, are three separate observed behaviours for which a single common model, or *grammar*, is to be inferred, then they may be fed to ATOM as the single string, /AB/CAA/ABC/, with / specified as a delimiter. This technique has been used extensively in grammatical inference with ATOM and is discussed in Ref. 90.

Certain descriptors may also be specified to be *inputs* to ATOM and then the surprise at their occurrence is not added in to the measure of approximation, so that neither *NE* nor *LE* take the predictions of these descriptors into account when evaluating the models. Conversely an input/output distinction may be inferred from the surprise patterns of ATOM's models in which descriptors giving rise to zero surprise are taken to be *outputs*, the others being inputs. This form of inference is illustrated in Ref. 89 where sample of the input/output behaviour of a deterministic, non-autonomous automaton is fed to ATOM without the input/output distinction being specified and this is readily inferred from the admissible models produced.

These computational studies illustrate clearly that the concepts developed in this paper are operational and can be applied to give clear analyses of actual data, to perform an automatic transformation from behaviour to model and thus identify the system giving rise to the behaviour.

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

5 DISCUSSION

In section 1 I mentioned a wide range of philosophical studies relevant to the problem of system identification. In this final section I shall link some of the concepts and results of these studies to specific aspects of the approach to identification developed in previous sections and exhibited in ATOM. These links are intended to be suggestive and analogical rather than a formal correspondence. There is an important sense in which the formal, mathematical approach to identification stands alone independent of its interpretation—a single phenomenon within it may assume many different roles under different interpretations and may be consistent with a variety of conflicting philosophical positions. Conversely there is an equally important sense in which each philosophical position stands per se and is wider, and more enduring, than an approach to identification couched in current mathematical and system-theoretic concepts.

System theorists are naturally wary of tangling with philosophy and the philosophical literature. System studies have generally arisen from two main sources: applications of technology, particularly automation and computers, on the one hand; and biological and environmental modelling on the other. These practical foundations are somewhat remote from the fundamental questions of philosophy, questions that to a large extent are dynamic and foundational just because no reasonable person would ask them. As Descartes remarks, "There is nothing so absurd or incredible that it has not been asserted by one philosopher or another". It is tempting to join Mach in his wistful remark, "I am a scientist and not a philosopher", but one also remembers that it is in his role of questioning established beliefs, as a philosopher of science, that his work had most impact on later developments such as relativity theory.[25]

Conversely, philosophers are wary of becoming too enmeshed in the science and mathematics of their time. Particularly in this epoch we have come to accept change as the natural order of things, and to assume that the scientific thought of today will not be that of tomorrow. *Nonstandard* as a technical term in logic and mathematics has also become something of a laudatory adjective. Pluralism, relativism and a willingness to take a legalistic, rather than theological, approach to scientific debate, arguing from axioms rather than beliefs, are the bases of modern scientific thought. They do not give the appearance of a base on which to lay firm foundations, unless laissez faire is itself seen as a metaphysical position.

These disclaimers aside, however, there is much to be gained by bringing to bear every possible approach, every insight of philsophy, system theory and practice, upon the fundamentally intractable problem of knowledge acquisition, of system identification. In the next section I shall discuss the approach developed in this paper in the context of Klir's hierarchical model of system epistemology, and then discuss both the approach and Klir's model in the context of other philosophical studies.

5.1 *Klir's Epistemological Hierarchy*

Figure 4 shows Klir's[104] hierarchy of epistemological levels of systems. The lowest level is one of *source systems*, effectively one of data definition whereby the way in which behaviour will be described is defined and agreed. In terms of section 2.2 of this paper, this is the level of definition of B, the set of possible behaviours. More specifically, in terms of section 3, it is the level of definition of D and of possible $b \in B$ mapping from some event space E to D, i.e. of possible **D-sets.** Thus level zero in the hierarchy might be said to define a *domain of discourse* for the description of behaviour. In terms of the example of section 4.2, a statement at this zero level domain of discourse is any string of A's and B's.

The next level in the hierarchy is one of *data systems*, effectively one of system observation whereby the actual behaviour of some system is described in terms of the agreed domain of discourse at level zero. In terms of this paper, this is the level of

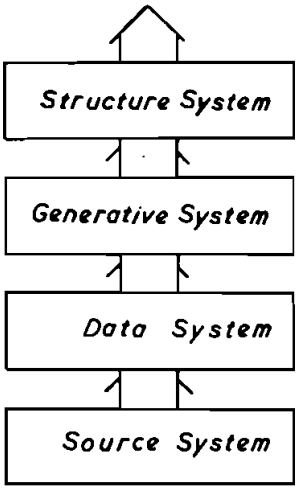

FIGURE 4 Klir's epistomological hierarchy.

definition of actual members of B, of a **D-set** representing behaviour. Thus level one of the hierarchy might be said to define the *observed* (or required) behaviour of the system. In terms of the example of section 4.2, the data system is that string of A's and B's shown in Figure 1a.

The next level of the hierarchy is one of *generative systems*, effectively one of a stationary model for a system. In the context of system identification the model will arise as a hypothesis about the generation of a behaviour described at level one. Klir notes that the model may be deterministic or stochastic and that its invariance (or stationarity) is with respect to a set of variables that may include space and time. Thus the uniform applicability of these concepts to both sequential behaviour and spatial patterns, as noted in section 3, is intrinsic to Klir's presentation also. In terms of section 2.2 of this paper, level two of the hierarchy is one of the models from a set M appropriate to a behaviour at level one. This paper may now be seen as concerned with the relationships between levels one and two in Klir's hierarchy. In terms of the example of section 4.2, any of the automata shown in Figure 1d is a generative system at level two.

The next level in the hierarchy is one of *structure systems*, effectively one at which the models themselves are seen to have internal structure and hence to be analysable in other terms. This level may be regarded as one at which a number of atomic models are taken as co-existent, with the system being described in terms of a relationship between them. Alternatively, it may be regarded as one of non-stationary models that are changing along some dimension such as space or time. There is no formal discussion of structure at level three in this paper. However, in section 4.2 I have noted informally that the 5-state model produced may be viewed as a 2-state system loosely coupled to a 3-state system, and this is a statement at the level three domain of discourse of Klir's hierarchy.

The hierarchy proceeds to *meta-levels*, etc., but the four lowest levels are most relevant to this paper. In particular the clear separation of the source system from the data system, and of the generative system from the structure system, are very helpful in defining the boundaries of the current study. I have defined the source level in very general terms as a class of **D-sets** and have not been concerned with properties, or appropriateness, at this level. I have defined the set of models also in very general terms and have been concerned only with the behaviour they generate, not their internal structures. The focus on the interface between levels one and two has been quite deliberate because of the previous lack of a sufficiently precise formulation of how this relationship should be established and evaluated. This has previously made the rigorous formulation of the problem of *system identification* (behaviour →model transformation) impossible, certainly for uncertain, e.g. stochastic or non-determinate, systems. The formulation in terms of approximation and complexity in this paper, and the specific definitions of approximation and complexity for monoids, etc., resolve this problem.

Technically, and practically, the study must be extended to levels zero and three to be of real value. The general **D-set** with its total lack of relations between data terms is unrealistic for many practical problems. We have topological or metrical constraints upon the semantics of our data terms which should be taken into account at higher levels. For example, a measurement of 2V at a control system transducer is *between* one of 1V and one of 5V; it may even be said to be *nearer* the 1V reading than the 5V one; it might even make sense to talk of its being *twice* the 1V value; and so on. Thus real measurements are rarely the arbitrary, unrelated symbols that are all that are hypothesized in the definition of a **D-set** used in this paper.

However, it should be noted that any restriction placed on B, or **D-sets**, reduces the ontological neutrality of our definition. Such a restriction reflects a preconception about possible behaviours and hence about possible structures underlying them, e.g. we may condition ourselves to observe only causal, or even only linear, systems. The trade-off, as illustrated in section 4.3, is between restrictions and speed-of-modelling/data-requirements. A preconception may prevent us from discovering a property of the data, but lack of it, when it is justified, may both increase our computational burden and lead us to less precise conclusions than are justified. We cannot *afford* to re-discover that a system is linear every time we wish to check its transfer function, but if we measure only the linear describing function we shall never realise that a nonlinearity has appeared.

The determination of the structure at level three of the models at level two is an important problem in its own right. In deterministic automata theory the Hartmanis–Stearns[191] and Krohn–Rhodes[192] decomposition techniques lead essentially to structural representations at level three of models at level two. Similar techniques may be applied to the probabilistic automaton models of section 4.2 to

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

give a formal basis for the statement that the best 5-state model corresponds to coupled 2-state and 3-state models. The hierarchical models of systems developed and analysed over a long period by Mesarovic[193–195] may be similarly automatically derived from model structures at level two, and much of "system theory" is concerned with such structural transformations.[196] There is the same inherent arbitrariness in transformations between levels two and three that I have already noted between levels one and two. There is no one structural decomposition of a model that is "correct" but many of varying preferability. In recent years this relationship has begun to be analysed also in terms of (structural) *complexity*.[197,198]

There is a direct relationship between measures of structural complexity at level three and the *model complexity* at level two assumed in this paper. The ordering of complexity over models required in section 2.2 corresponds to a level three construction of the least-complex structure for the model at level two, e.g. a *minimum-state* structure. Thus the counting of states that is taken as a basis for evaluating complexity in the specific automata-theoretic studies of this paper corresponds to a particular concept of the nature of systems at level three. However, it is only one possible measure even in the context of automata—we may well feel than an 8-state automaton that can be regarded as a certain form of product of three 2-state automata is rather less complex than one which cannot. This is one reason for defining complexity to be only a *partial ordering* over models since this type of consideration does not necessarily lead to a simple numeric measure of complexity giving a chain order.

Such considerations are clearly of great importance but they do not affect the discussion of this paper where no constraint is placed on the ordering of model complexity, except to reinforce the statement in section 2.1 that the order relation is *not* intrinsic to the class of models. We do not normally envisage a model without also assuming a structure but, in fact, the same model may be regarded as having many different structures and complexities.

5.2 *The Wider Context*

The wider philosophical context for Klir's epistemological hierarchy and the operation of ATOM is far too wide to cover in this paper, It ranges through the whole gamut of philosophical positions and concepts mentioned in section 1, and does so necessarily in the sense that focussing on only one aspect itself distorts the overall picture. Much of the relevant discussion has taken place in the context of the behavioural sciences,[199,200] including biology,[201–203] psychology,[204–209] sociology,[210–214] and economics,[116,146,147] because the problem of identification in the technological sciences is so much more readily defined—we are already aware of the structure, or structural possibilities, of systems that we ourselves have designed. The frontiers of the physical sciences stand between these extremes and have aspects of both,[215–218] from the deep epistemological problems of particle physics[219–220] and time,[221–223] to the logical, postulational systems of dynamics.[224,225] All of these differing subject areas have their own, varied but related, approaches to the problem of knowledge acquisition, and the problem of system identification, of behaviour →structure inference, is common to them all.

As noted in section 1, the formal approach to the inference of system properties developed in this paper, and by others,[79–82,104] avoids rather than resolves many difficult problems. Hume's inductive scepticism leaves us with no logical justification of our processes of knowledge acquisition and hence what we actually do becomes somewhat arbitrary. This arbitrariness shows up in the present study at a number of levels: the choice of **D-sets** and their structures (if any); the choice of models; the ordering relation of complexity; the induced order relations of approximation; and so on. The effect of additional hypotheses to resolve some of the arbitrary decisions necessary has been discussed, and can be taken into account formally once the *decision* has been taken to include them.

Dilman (Ref. 34, p. 19) argues that, although such decisions are *groundless*, they are, "not arbitrary in the sense that so much in our life hangs together with them." Wittgenstein emphasizes this view many times stating that, "Scepticism is not irrefutable, but obviously nonsensical when it tries to raise doubts where no question can be asked" (Ref. 226, 6.51), and, "what the law of causality is meant to exclude cannot even be described" (Ref. 226, 6.362). It is the basis of his remark,[33] "If anyone said that information about the past could not convince him that something would happen in the future, I should not understand him." This line of argument establishes a direct link between Simon's[79] "normative theory of law discovery" and our immediate, predictive use of the "laws," even though they are only "patterns in data." The historian may be

interested in models as providing a rational explanation of past events, but usually we intend them as guides to the future.

None of this, or any variant of the counter-arguments to Hume,[35-40] actually refutes the need for arbitrary (or, at least, groundless) decisions, and, as a rejoinder to one line of argument that occurs at a number of levels, it is worth noting the meta-theoretic result that: *the introduction of any hypothesis into the problem of identification allows the generation of counter-examples that do not satisfy the hypothesis and are modelled in a meaningless way*. This is rather too vague for formal proof at present but there are now sufficient examples to make it a reasonable inductive assertion. We cannot expect to introduce useful constraints on the inference process that are also ontologically neutral, and if we make *any* ontological commitment then we are in danger of not comprehending what are actually simple phenomena.

This is one of the problems, with its connotations of arbitrary decisions and commitment, that has lead to so-called *existentialist* philosophical studies, although the diversity of positions taken is such that one cannot speak of an existentialist "school." These studies may appear too introspective, often verging on solipsism and nihilism, to those whose standard is the classical philosophy of science. However, there is much to be gained in the understanding of the foundations of system theory, particularly in its application to the behavioural sciences, through the works of, for example, Merleau-Ponty,[227] Sartre,[228] Heidegger[7] and Husserl.[16] Perhaps the motivation for studying the questions thus raised best comes from Becker's rather more humorous Yale lectures on the, "heavenly city of eighteenth-century philosophers."[229] To laugh at the preconceptions of the past in successive historic progression generates the momentum necessary to carry history into the future and allow us, as if retrospectively, to view the preconceptions of the present.

The authors noted above are associated with the *phenomenological* approach to the study of knowledge, whose proponents again range from the mystical to the scientifically respectable, notably Mach.[25,230] Phenomenology is concerned with the study of experience with a view to bringing out their "essences," their underlying "reason." For Husserl,[16,231] it was to be an analysis "free from presuppositions," as contrasted to the *a priori* bases for most metaphysics. Thus phenomenological analysis is concerned initially with the lowest level of Klir's hierarchy, one that is neglected, to some extent unrecognized, in classical systems theory. We are used to the terms in which behaviour is to be described being specified for us in advance and tend to regard the problem of identification as commencing beyond this point—a general method should be able to cope with any domain of discourse at a lower level. However, the terms of reference accepted at level zero propagate upwards through the hierarchy and dominate all higher-level concepts. The nature of such effects is best seen through the phenomenological case histories discussed by Spiegelberg[232] (e.g. the phenomenology of "force"), and Roche.[233] This last is of particular interest also for its comparison of phenomenological analysis with the "conceptual analysis" of logical positivism.[234]

Once above the level of existence and the phenomenon, the links between the problem of identification, Klir's hierarchy, and philosophical discussions are more obvious. The *structuralism*[235] and *presuppositional*[236] analyses of modern linguistic philosophy and the linguistic orientation of logical positivism clearly have their places in what may be regarded as grammatically orientated inductive inference process. It is easy to integrate this approach to the problem of identification with what Giedymin[237] terms the six doctrines of "strict positivism." However, equally there is no need to accept these doctrines—the order relations of section 2.2 define a form of *rationality*, and hence the basis for a "rationality principle,"[6,238] but this need be none that we know or accept at present.

Even with all the low level choices made, with source level, data level, and the class of models at generative level together with their orderings, well determined, there are still decisions to be made. For example, the regular events corresponding to each of the 1 through 5-state admissible models of the behaviour analysed by ATOM in section 4.2 are each viable hypotheses as to the structure of the data. It was assumed in that section that we recognize the 5-state model as that which is "correct," but it may well be that the improvement in approximation from 1 to 2 states is enough to satisfy us—"facts are always presented at a sacrifice of completeness and never with greater precision than fits the needs of the moment" (Mach, Ref. 25, p. 174). Having hypothesized that the event structure is of the form (BB*A)*, however, how do we proceed? Are we, with Popper,[41] to search for the sequence AA that refutes our hypothesis or, with Putman,[44] to search for the sequence BBB that is predicted but not yet seen. In ATOM terms, finding the first would force us back to the 1-state model,

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

whereas not finding the second would push us on to the 3-state model. There is an interplay not only between the consistency of data and theory but also between the complexity of a theory and its value in approximating the data—Newtonian mechanics and special relativity are both *admissible* theories on all current data in the formal sense defined in this paper.

Finally, the discussion of the preceding section serves to illustrate the role of programs such as ATOM as themselves sources of phenomena relating to inductive inference and knowledge acquisition. The *gedanken experiment* of having a *confirmation machine*[239] is itself a source of stimulating discussion, but it is better by far to actually *have* one and to see it in action, both on data for which one "knows the answer" and on data for which one does not. There is nowadays no reason why any theory of knowledge acquisition should not be put into operational form and allowed, or required, to compete in real and artificial worlds. If we cannot understand *automata* then how can we hope to comprehend *ourselves*, and if we cannot make sense of what we *can* do then how may we say what we *should* do?

6 SUMMARY AND CONCLUSIONS

The problem of system identification has effectively been defined as: "given a sample of the behaviour of some system, to determine the admissible subspace of some prescribed class of models that would account for it." Admissibility is itself defined in terms of two order relations on models: a static order of complexity; and a dynamic order of approximation induced by the behaviour. An admissible model is such that any other with better approximation has greater complexity. It has been shown that this framework encompasses the identification of both deterministic and stochastic finite-state automata, or grammars, but that it also encompasses non-temporal "behaviour" in, for example, pattern recognition.

The behaviour of the plot of approximation against complexity for admissible models has been investigated. It has been shown (section 3.1) that for zero approximation (perfect fit) nearly all behaviours require models of nearly maximum complexity. A general basis for the measurement of approximation has been introduced related to work on logical, or semantic, probability and the elicitation of subjective probabilities (section 3.2), and it is shown that, with probabilistic sources, the approximation of an admissible model may be regarded as an entropy measure of the behaviour (section 3.3).

The role of structural constraints upon models has been examined and concepts of causality, locality and uniformity have been introduced as "rules of the game" (section 3.3.1). The way in which the placement of *state* "tokens" allows local algorithms to recognize global properties has been discussed together with the concept of observability. Finally, the sense in which an identification algorithm can determine the structure of a probabilistic automaton has been analysed (section 3.3.2).

A computer implementation of an algorithm for determining admissible models of sequential behaviour, ATOM, has been described (section 4.1) and its operation illustrated through the analysis of some sample behaviours, including coupled automata (section 4.2) and non-recurrent sequences (section 4.3). Plots of the approximation against complexity, admissible models, and the "surprise" of the models at each description in the behaviour have been given. In particular, the capability of ATOM to infer, or specify in advance, special properties of the descriptors, such as being delimiters or inputs has been discussed and illustrated (section 4.4).

In the discussion section the behaviour of ATOM and the general approach to identification proposed have been analysed in terms of Klir's epistemological hierarchy of systems. It has been shown that the problem solved is one of the interface between levels one and two, of data systems and generative systems (section 5.1). The relationship of these to both lower and higher levels, and in particular the relationship of the model complexity used to structural complexity at level two have also been discussed. Finally, a brief summary has been presented of the relationship of these studies to various philosophical developments (section 5.2).

This paper gives a precise and complete basis for the analysis and solution of one of the major problems of system identification, that of behaviour →model transformation, effectively the movement from level one to level two of Klir's hierarchy. In that sense it "solves" the technical problem of structural, or grammatical inference, for nondeterministic and stochastic automata, and for a wide class of similar problems. The practicality of the solution is limited by the computational effort required, but it may well be that when we find ourselves searching large numbers of models we

should ask whether we have the right formulation—would another model space be more appropriate?; and whether we are solving the right problem—would a sub-optimal solution be equally acceptable? Both questions are analysable within the framework of this paper.

On the other hand, I have attempted to indicate the way in which a formal solution of the interrelation of levels one and two relates to consideration at level zero, the *phenomenological* level, and level three, the *structural* level, and how all these levels relate to the various approaches to knowledge acquisition studied in various philosophical contexts. The availability nowadays of powerful interactive computers makes it possible to implement the system-theoretic methodologies imbedded in these philosophical positions and investigate them not as theories but as practice. It is through observation of, and interaction with, the autonomous operation of our theories, as computer algorithms competing with one another in the acquisition of knowledge of real and artificial data, that we have most to gain at present.

Downloaded by [University of Calgary] at 10:52 14 February 2012

ACKNOWLEDGEMENTS

Many people have influenced the course of development of this paper and the work reported in it. John Andreae, now at the University of Canterbury, New Zealand, and Alan Watson, at the Psychological Laboratory, Cambridge, did much to motivate the problem through wide-ranging discussions and their own approaches. ATOM itself developed in a competitive atmosphere in my own laboratory with the stimulation of colleagues, notably Peter Facey and Ian Witten. I am grateful to Michael Arbib, Richard Dawkins, Joe Goguen, George Klir, Ladislav Kohout, Steve Matheson and Roger Moore, for helpful discussions and comments, and to Judea Pearl of UCLA for detailed criticism of the work. Finally, thanks to my wife (who always complains she is left out at this stage) for not clearing up the mess that, for some weeks, was this paper and its references.

REFERENCES

1. J. Locke, *An Essay Concerning Human Understanding.* (Ed. J. W. Yolton), Dent, London, 1961 (original 1690).
2. I. Kant, *Critique of Pure Reason.* (Trans. J. M. D. Meiklejohn), George Bell, London, 1897 (original 1781).
3. G. Berkeley, *The Principles of Human Knowledge.* (Ed. G. J. Warnock), Fontana, London, 1962 (original 1710).
4. Plato, *The Republic.* (Trans. H. D. P. Lee), Penguin, Middlesex, 1955 (original cf. 380 BC).
5. D. Hume, *Enquiries Concerning Human Understanding.* (Ed. L. A. Selby-Bigge, revised P. H. Nidditch), Clarendon Press, Oxford, 1975 (original 1777).
6. K. R. Popper, *Objective Knowledge.* Clarendon Press, Oxford, 1972.
7. M. Heidegger, *Existence and Being.* Vision, London, 1949.
8. R. Carnap, *Logical Foundations of Probability.* University of Chicago Press, Chicago, 1950.
9. R. Carnap, *The Continuum of Inductive Methods.* Chicago University Press, Chicago, 1952.
10. P. A. Schilpp (ed.), *The Philosophy of Rudolph Carnap.* Open Court Press, La Salle, Illinois, 1963.
11. R. C. Buck and R. S. Cohen (eds.), *PSA 1970: In Memory of Rudolph Carnap.* D. Reidel, Dordrecht, Holland, 1971.
12. R. Carnap and R. C. Jeffrey (eds.), *Studies in Inductive Logic and Probability.* University of California Press, Berkeley, 1971.
13. R. Swinburne, *An Introduction to Confirmation Theory.* Methuen, London, 1973.
14. H. J. Blackhan, *Six Existentialist Thinkers.* Routledge and Kegan Paul, London, 1952.
15. F. Brentano, *Psychology from an Empirical Standpoint.* (Ed. L. L. McAlister, trans. D. B. Terrell, A. C. Rancurello, and L. L. McAlister), London, 1973.
16. E. Husserl, *Phenomenology and the Crisis in Philosophy.* (Trans. Q. Lauer) Harper, New York, 1965.
17. E. Pivcevic (ed.), *Phenomenology and Philosophical Understanding.* Cambridge University Press, Cambridge, 1975.
18. N. R. Hanson, *Patterns of Discovery.* Cambridge University Press, 1958.
19. M. Polanyi, *Personal Knowledge.* Harper and Row, New York, 1964.
20. C. R. Kordig, "The theory-ladeness of observation." *Reviews of Metaphysics*, **24**, 1971, pp. 448–484.
21. G. Gale and E. Walter, "Kordig and the theory-ladeness of observation." *Philos. Sci.*, **40**, 1973, pp. 415–432.
22. M. B. Hesse, *Models and Analogies in Science.* University of Notre Dame Press, Indiana, 1966.
23. D. C. Dorrough, "A logical calculus of analogy involving functions of order 2." *Notre Dame Journal of Formal Logic*, **11**, 1970, pp. 321–336.
24. A. I. Uemov, "The basic forms and rules of inference by analogy." In: P. V. Tanasec (ed.), *Problems of the Logic of Scientific Knowledge.* D. Reidel, Dordrecht, 1970, pp. 266–311.
25. J. T. Blackmore, *Ernst Mach.* University of California Press, Berkeley, 1972, pp. 173–174.
26. H. R. Post, "Simplicity and scientific theories." *Brit. J. Philos. Sci.*, **11**, 1960, pp. 32–41.
27. E. Sober, *Simplicity.* Clarendon Press, Oxford, 1975.
28. D. K. Lewis, *Convention: A Philosophical Study.* Harvard University Press, Cambridge, Mass., 1969.
29. M. A. E. Dummett, "The justification of deduction." *Proc. Brit. Acad.*, **59**, 1973, pp. 1–34.
30. S. Haack, "The justification of deduction." *Mind*, **85**, 1976, pp. 112–119.
31. E. H. Madden, "Hume and the fiery furnace." *Philos. Sci.*, **38**, 1971, pp. 64–78.
32. D. C. Stove, *Probability and Hume's Inductive Scepticism.* Clarendon Press, Oxford, 1973.
33. L. Wittgenstein, *Philosophical Investigations* (trans. G. E. M. Anscombe). Basil Blackwell, Oxford, 1974, section 481 (original 1945).
34. I. Dilman, *Induction and Deduction.* Basil Blackwell, Oxford, 1973.
35. H. Reichenbach, *The Theory of Probability.* University of California Press, Berkeley, 1949.
36. R. Harrod, *Foundations of Inductive Logic.* Macmillan, London, 1956
37. J. J. Katz, *The Problem of Induction and Its Solution.* University of Chicago Press, Chicago, 1962.

38. M. Black, *Margins of Precision.* Cornell University Press, Ithaca, 1970.
39. N. Rescher, *The Primacy of Practice.* Basil Blackwell, Oxford, 1973.
40. R. Swinburne (ed.), *The Justification of Induction.* Oxford University Press, 1974.
41. K. R. Popper, *The Logic of Scientific Discovery.* Hutchinson, London, 1959.
42. K. R. Popper, *Conjectures and Refutations.* Routledge and Kegan Paul, London, 1963.
43. P. A. Schilpp (ed.), *The Philosophy of Karl Popper.* Open Court Press, La Salle, Illinois, 1974.
44. H. Putnam, "The 'corroboration' of theories." In Ref. 43 and in H. Putnam, *Mathematics, Matter and Method.* Cambridge University Press, 1975, pp. 250–269.
45. A. C. Michalos, *The Popper-Carnap Controversy.* Martinus Nijhoff. The Hague, 1971.
46. T. S. Kuhn, *The Structure of Scientific Revolutions.* University of Chicago Press, Chicago, 1962.
47. R. K. Merton, *The Sociology of Science.* (Ed. N. W. Storer), University of Chicago Press, 1973.
48. I. Lakatos, "Falsification and the methodology of scientific research programmes." In: I. Lakatos and A. Musgrave (eds.) *Criticism and the Growth of Knowledge.* Cambridge University Press, Cambridge, 1970.
49. M. Hesse, *The Structure of Scientific Inference.* Macmillan, London, 1974.
50. P. Feyerabend, *Against Method.* NLB, London, 1975.
51. E. Gellner, *Legitimation of Belief.* Cambridge University Press, Cambridge, 1974.
52. I. Hacking, *The Emergence of Probability.* Cambridge University Press, 1975.
53. L. J. Savage, *Foundations of Statistics.* John Wiley, New York, 1954.
54. M. H. Foster and M. L. Martin (eds.) *Probability, Confirmation, and Simplicity.* Odyssey Press, New York, 1966.
55. H. E. Kyburg, *Probability and Inductive Logic.* Macmillan, London, 1970.
56. G. I. Ruzavin, "Probability logic and its role in scientific research." In: P. V. Tanavec (ed.), *Problems of the Logic of Scientific Knowledge.* D. Reidel, Dordrecht, Holland, 1970, pp. 212–265.
57. T. L. Fine, *Theories of Probability.* Academic Press, New York, 1973.
58. B. de Finetti, *Probability, Induction and Statistics.* John Wiley, London, 1972.
59. G. Maxwell and R. M. Anderson (eds.), *Induction, Probability and Confirmation.* University of Minnesota Press, Minn., 1975.
60. Y. Bar-Hillel and R. Carnap, "Semantic information." *Brit. J. Philos. Sci.*, **4**, 1953, pp. 147–157.
61. Y. Bar-Hillel, *Language and Information.* Addison-Wesley, Reading, Mass., 1964.
62. J. Hintikka, "On semantic information." In J. Hintikka and P. Suppes (eds.), *Information and Inference.* D. Reidel, Dordrecht, 1970, pp. 3–27.
63. R. Hilpinen, "On the information provided by observations." Ibid. pp. 97–122.
64. R. C. Jeffrey, *The Logic of Decision.* New York, 1965.
65. I. Levi, *Gambling with Truth.* M.I.I. Press, Mass., 1967.
66. G. Menges (ed.), *Information, Inference and Decision.* D. Reidel, Dordrecht, Holland, 1974.
67. P. Martin-Löf, "The definition of random sequences." *Inf. Control*, **9**, 1966, pp. 602–619.
68. A. Kolmogorov, "Logical basis for information theory and probability theory." *IEEE Trans. Inf. Theory*, **IT-14**, 1968, pp. 662–664.
69. D. G. Willis, "Computational complexity and probability constructions." *J. Assn. Comput. Mach.*, **17**, 1970, pp. 241–259.
70. C. P. Schnorr, "A unified approach to the definition of random sequences." *Math. Syst. Theory*, **5**, 1971, pp. 246–258.
71. G. J. Chaitin, "A theory of program size formally identical to information theory." *J. Assn. Comput. Mach.*, **22**, 1975, pp. 329–340.
72. A. Lempel and J. Ziv, "On the complexity of finite sequences." *IEEE Trans. Inf. Theory*, **IT-22**, 1976, pp. 75–81.
73. J. A. Feldman, "Some decidability results on grammatical inference and complexity." *Inf. and Control*, **20**, 1972, pp. 244–262.
74. K. S. Fu and T. L. Booth, "Grammatical inference: Introduction and survey—Part II." *IEEE Trans. Syst., Man and Cyber.*, **SMC-5**, 1975, pp. 409–423.
75. A. R. Patel, *Grammatical Inference for Probabilistic Finite-State Languages*, Ph.D. thesis, University of Connecticut, USA, 1972.
76. F. J. Maryanski, *Inference of Probabilistic Grammars.* Ph.D. thesis, Dep. Elec. Eng. and Comput. Sci., Univ. Conn., Storrs, 1974.
77. H. Putnam, "Robots: Machines or artificially created life." *J. Philos.*, **61**, 1964, pp. 668–691.
78. R. E. Kling, "A paradigm for reasoning by analogy." *Proc. 2nd Int. Joint Conf. in Art. Intell.* British Computer Society, London, 1971, pp. 568–585.
79. H. A. Simon, "Does scientific discovery have a logic?" *Philos. Sci.*, **40**, 1973, pp. 471–480.
80. A. Newell and H. A. Simon, *Human Problem Solving.* Prentice-Hall, New Jersey, 1972.
81. R. J. Solomonoff, "A formal theory of inductive inference." *Information and Control*, **7**, 1964, pp. 1–22, 224–254.
82. G. J. Klir, "On the representation of activity arrays." *Int. J. General Systems*, **2**, 1975, pp. 149–168.
83. R. J. Nelson, "On machine expectation." *Synthese*, **31**, 1975, pp. 129–139.
84. J. H. Andreae and J. G. Cleary, "A new mechanism for a brain." *Int. J. Man-Machine Studies*, **8**, 1976, 89–119.
85. L. A. Zadeh, "From circuit theory to system theory." *Proc. IRE*, **50**, 1962, pp. 856–865.
86. P. Eykhoff, *System Identification.* John Wiley, London, 1974.
87. R. M. Wharton, "Approximation language identification." *Information and Control*, **26**, 1974, pp. 236–255.
88. B. R. Gaines, "Approximate identification of automata." *Electron Lett.*, **11**, 1975, pp. 444–445.
89. B. R. Gaines, "Behaviour/structure transformations under uncertainty." *Int. J. Man-Machine Studies*, **8**, 1976, pp. 337–365.
90. B. R. Gaines, "Inference of Stochastic grammars: a formulation and solution." EES-MMS-AUT-76, Department of Electrical Engineering Science, University of Essex, Colchester, Essex, 1976.
91. C. A. B. Smith, "Consistency in statistical inference and decision." *J. Roy. Statistic Soc.*, ser. B, **23**, 1961, pp. 1–37.
92. C. A. B. Smith, "Personal probability and statistical analysis." *J. Roy Statist. Soc.*, ser. A, **128**, 1965, pp. 469–499.
93. J. Aczel and J. Pfanzagl, "Remarks on the measurement of subjective probability and information." *Metrika*, **2**, 1966, pp. 91–105.

Downloaded by [University of Calgary] at 10:52 14 February 2012

94. E. H. Shuford, A. Albert and H. E. Massengill, "Admissible probability measurement procedures." *Psychometrika*, **31**, 1966, pp. 125–145.
95. R. L. Winkler and A. H. Murphy, "Good probability assessors." *J. App. Meteorology*, **7**, 1968, pp. 751–758.
96. L. J. Savage, "Elicitation of personal probabilities and expectations." *J. Amer. Stat. Assn.*, **66**, 1971, pp. 783–801.
97. R. L. Winkler, "Probabilistic prediction: some experimental results." *J. Amer. Stat. Assn.*, **66**, 1974, pp. 625–688.
98. R. M. Hogarth, "Cognitive processes and the assessment of subjective probability distributions." *J. Amer. Stat. Assn.*, **70**, 1975, pp. 271–294.
99. E. Shuford and T. A. Brown, "Elicitation of personal probabilities and their assessment." *Instructional Sci.*, **4**, 1975, pp. 137–188.
100. J. Pearl, "An economic basis for certain methods of evaluating probabilistic forecasts." UCLA-ENG-7561, School of Engineering and Applied Science, UCLA, Calif., July 1975.
101. J. Pearl, "On the complexity of inexact computations." UCLA-ENG-0775, ibid, July 1975.
102. J. Pearl, "On the complexity of imprecise causal models." UCLA-ENG-7560, ibid., July 1975.
103. J. Pearl, "On the complexity of computing probabilistic assertions." UCLA-ENG-7562, ibid., July 1975.
104. G. J. Klir, "Identification of generative structures in empirical data." *Int. J. General Systems*, **3**, 1976, pp. 89–104.
105. G. J. Klir, "Processing of fuzzy activities of neutral systems." *In: Progress in Cybernetics and Systems Research.* (R. Trappl and F. R. Pichler, eds.), Hemisphere, Washington, D.C., 1975, pp. 21–24.
106. J. A. Goguen, "Realization is universal." *Mathematical Systems Theory*, **6**, 1973, pp. 359–374.
107. J. A. Goguen, "Discrete-time machines in closed monoidal categories." *Int. J. Comput. Syst. Sci.*, **10**, 1975, pp. 1–43.
108. L. S. Bobrow and M. A. Arbib, *Discrete Mathematics.* W. B. Saunders, Philadelphia, 1974.
109. M. A. Arbib and E. G. Manes, "Foundations of system theory: decomposable systems." *Automatica*, **10**, 1974, pp. 285–302.
110. H. Ehrig and H. J. Kreowski, "Systematic approach of reduction and minimization in automata and system theory." Report Nr. 73-16, Technische Universität Berlin, Oct. 1973.
111. H. Ehrig, *Universal Theory of Automata.* B. G. Teubner, Stuttgart, 1974.
112. B. R. Gaines, "The role of randomness in system theory." EES-MMS-RAN-76, Department of Electrical Engineering, University of Essex, Colchester, U.K., January 1976.
113. N. Rescher, *Scientific Explanation.* Free Press, New York, 1970.
114. P. Suppes, *Probabilistic Metaphysics.* Filofiska Studier, Uppsala University, Sweden, 1974.
115. J. L. Mackie, *The Cement of the Universe.* Clarendon Press, Oxford, 1975.
116. P. D. McClelland, *Causal Explanation and Model Building.* Cornell University Press, 1975.
117. B. R. Gaines, "Foundations of fuzzy reasoning." *Int. J. of Man-Machine Studies*, **8**, 1976, to appear.
118. B. R. Gaines, "On the complexity of causal models." *IEEE Trans. on Systems, Man and Cybernetics*, **SMC-6**, 1976, pp. 56–59.
119. L. Weiss, *Statistical Decision Theory.* McGraw-Hill, New York, 1961.
120. H. Kwakernaak, "Admissible adaptive control." In: *Proc. IFAC Symposium on The Theory of Self-Adaptive Control Systems*, Society of Instrument Technology, London, 1965.
121. B. R. Gaines and I. H. Witten, "Stability and admissibility of adaptive threshold logic convergence." *IEEE Trans. on Computers*, **C-25**, 1976, to appear.
122. A. Nerode, "Linear automata transformations." *Proc. Amer. Math. Soc.*, **9**, 1958, pp. 541–544.
123. M. A. Arbib, *Theories of Abstract Automata.* Prentice-Hall, New Jersey, 1969.
124. M. O. Rabin and D. Scott, "Finite automata and their decision problems." *IBM J. Res. Dev.*, **3**, 1959, pp. 114–125.
125. J. Hopcroft, "An $n \log n$ algorithm for minimizing states in a finite automaton." In: Z. Kohavi and A. Paz (eds.), *Theory of Machines and Computations.* Academic Press, New York, 1971, pp. 189–196.
126. A. M. Mathai and P. N. Rathie, *Basic Concepts in Information Theory and Statistics.* Wiley Eastern, New Delhi, 1975.
127. J. A. Goguen, "Concept representation in natural and artificial languages: axioms, extensions and applications for fuzzy sets." *Int. J. Man-Machine Studies*, **6**, 1974, pp. 513–561.
128. B. R. Gaines and L. J. Kohout, "The logic of automata." *Int. J. General Systems*, **2**, 1975, pp. 191–208.
129. A. Kaufmann, *Theory of Fuzzy Subsets, Vol. 1.* (Trans. D. L. Swanson). Academic Press, New York, 1975.
130. L. A. Zadeh, K. S. Fu, K. Tanaka, and M. Shimura (eds.), *Fuzzy Sets and Their Applications to Cognitive and Decision Processes.* Academic Press, New York, 1975.
131. M. A. Arbib and E. G. Manes, "A category-theoretic approach to systems in a fuzzy world." *Synthese*, **30**, 1975, pp. 381–406.
132. L. A. Zadeh, "A fuzzy-algorithmic approach to the definition of complex or imprecise concepts." *Int. J. Man-Machine Studies*, **8**, 1976, pp. 249–291.
133. B. R. Gaines, "Fuzzy reasoning and the logics of uncertainty." EES-MMS-UNC-75 to appear in *Proc. 6th Annual Symposium on Multiple-Valued Logic*, Utah, May 1976.
134. B. R. Gaines and L. J. Kohout, "The fuzzy decade: a bibliography of fuzzy systems and closely related topics." *Int. J. Man-Machine Studies*, **8**, 1976, to appear.
135. P. H. Sellers, "An algorithm for the distance between two finite sequences." *J. Combinatorial Theory (A)*, **16**, 1974, pp. 253–258.
136. R. A. Wagner and M. J. Fischer, "The string-to-string correction procedure," *J. Assn. Comput. Mach*, **21**, 1974, pp. 168–173.
137. W. M. Fitch and E. Margolias, "Construction of phylogenetic trees." *Science*, **155**, 1967, pp. 279–284.
138. D. Sankoff, "Matching sequences under deletion insertion constraints." *Proc. Nat. Acad. Sci. U.S.A.*, **69**, 1972, pp. 4–6.
139. R. Carnap, "The aim of inductive logic." In: E. Nagel, P. Suppes and A. Tarski (eds.), *Logic, Methodology and Philosophy of Science*, Stanford University Press, 1962, pp. 303–318.
140. I. J. Good, "Subjective probability as the measure of a non-measurable set." Ibid., pp. 319–329.
141. G. H. von Wright, "Remarks on the epistemology of subjective probability." Ibid. pp. 330–339.
142. C. Villegas, "On qualitative probability σ-algebras." *Anns. Math. Stat.*, **35**, 1964, pp. 1787–1796.
143. J. M. Vickers, "Some remarks on coherence and subjective probability." *Philos. Sci.*, **32**, 1965, pp. 32–38.

Downloaded by [University of Calgary] at 10:52 14 February 2012

Downloaded by [University of Calgary] at 10:52 14 February 2012

144. G. Menges, "On subjective probability and related problems." *Theory and Decision*, **1**, 1970, pp. 40–60.
145. B. Grofman and G. Hyman, "Probability and logic in belief systems." *Theory and Decision*, **4**, 1973, pp. 179–195.
146. G. L. S. Shackle, *Uncertainty in Economics*. Cambridge University Press, 1955.
147. G. L. S. Shackle, *Decision, Order and Time*. Cambridge University Press, 1961.
148. N. Rescher, *Many-Valued Logic*. McGraw-Hill, 1969, pp. 184–188.
149. R. R. Kalman, "Nonlinear aspects of sampled-data control systems." In: *Proc. Symp. Nonlinear Circuit Analysis*, Polytechnic Institute of Brooklyn, 1957, pp. 273–313.
150. B. R. Gaines, "A calculus of possibility, eventuality and probability." In EES-MMS-FUZI-75, Department of Electrical Engineering, University of Essex, Colchester, U.K., 1975.
151. N. Rescher, "A probabilistic approach to model logic." *Acta Philosophica Fennica*, **16**, 1963, pp. 215–226.
152. S. Danielsson, "Modal logic based on probability theory." *Theoria*, **33**, 1967, pp. 189–197.
153. S. Miura, "Probabilistic models of modal logics." *Bull. Nagoya Inst. Tech.*, **24**, 1972, pp. 67–72.
154. G. E. Hughes and M. J. Creswell, *An Introduction to Modal Logic*. Methuen, London, 1968.
155. D. P. Snyder, *Modal Logic*, Van Nostrand, New York, 1971.
156. A. Prior, *Past, Present and Future*. Clarendon Press, Oxford, 1967.
157. J. Aczel, "On different characterizations of entropies." In: M. Behara, K. Krickeberg and J. Wolfowitz (eds.), *Probability and Information Theory*, Lecture Notes in Mathematics, **89**, Springer-Verlag, Berlin, 1969.
158. S. Watanabe, *Knowing and Guessing*. John Wiley, New York, 1969.
159. T. G. Windeknecht, "Mathematical Systems theory: causality." *Math. Syst. Theory*, **1**, 1967, pp. 279–288.
160. T. G. Windeknecht, *General Dynamical Processes*. Academic Press, New York, 1971.
161. P. J. Orava, "Notion of dynamical input-output systems: causality and state concepts." *Int. J. Syst. Sci.*, **5**, 1971, pp. 793–806.
162. D. J. H. Moore, "A theory of form." *Int. J. Man-Machine Studies*, **3**, 1971, pp. 31–59.
163. R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. John Wiley, New York, 1973.
164. W. C. Hoffman, "The Lie algebra of visual perception." *J. Math. Psychol.*, **3**, 1966, pp. 65–98.
165. U. Grenander, "Foundations of pattern analysis." *Quart. App. Math.*, **27**, 1969, pp. 1–55.
166. M. Minsky and S. Papert, *Perceptrons*. Massachusetts Institute of Technology, 1969.
167. M. E. Hellman and T. M. Cover, "Learning with finite memory." *Ann. Math. Statistics*, **41**, 1970, pp. 765–782.
168. I. H. Witten, "The apparent conflict between estimation and control—a survey of the two-armed bandit problem." *J. Franklin Inst.*, **301**, 1976, pp. 161–189.
169. M. Perles, M. O. Rabin and E. Shamir, "Theory of definite automata." *IEEE Trans. Electron. Comp.*, EC-12, 1963, pp. 233–243.
170. Y. Zalcstein, "Locally testable languages." Technical Report, Department of Computer Science, Carnegie-Mellon University, 1971.
171. R. McNaughton and S. Papert, *Counter-Free Automata*. Massachusetts Inst. Tech., 1971.
172. G. Birkhoff, *Dynamical Systems. Amer. Math. Soc. Coll. Pub.*, 1927.
173. A. A. Markov, *Theory of Algorithms*. Akad. Nauk SSSR, 1954 (trans. Israel Program for Scientific Translation, Jerusalem, 42, 1962).
174. A. C. de Fiorino, "Generalized Markov algorithms and automata." In E. R. Caianiello (ed.), *Automata Theory*, Academic Press, New York, 1966.
175. L. A. Zadeh, "The concept of state in system theory." In: M. D. Mesarovic (ed.), *Views on General Systems Theory*, John Wiley, New York, 1964, pp. 39–50.
176. L. A. Zadeh, "The concepts of system, aggregate, and state in system theory." In: L. A. Zadeh and E. Polak (eds.), *System Theory*, McGraw-Hill, 1969, pp. 3–42.
177. S. Salovaara, "On set theoretical foundations of system theory—a study of the state concept." *Acta. Polytech. Scand.*, No. 15, 1974, pp. 1–74.
178. R. L. Gregory, *The Intelligent Eye*. Weidenfeld and Nicolson, London, 1970.
179. C. A. J. Vleck, "Multiple probability learning." *Acta Psychologica*, **33**, pp. 207–232, 1970.
180. P. Suppes and W. Rottmayer, "Automata." In: E. C. Carterette and M. P. Friedman (eds.), *Handbook of Perception*, Vol. 1, *Historical and Philosophical Roots of Perception*, Academic Press, New York, 1974.
181. R. Dawkins and M. Dawkins, "Decisions and the uncertainty of behaviour." *Behaviour*, **45**, pp. 83–103, 1973.
182. M. Dawkins and R. Dawkins, "Some descriptive and explanatory stochastic models of decision-making." In: D. J. McFarland (ed.), *Motivational Control Systems Analysis*, Academic Press, 1974, pp. 119–168.
183. P. V. Facey and B. R. Gaines, "Real-time system design under an emulator embedded in a high-level language." *Proc. Brit. Comput. Soc.*, DATAFAIR 73, Apr. 1973, pp. 285–291.
184. B. R. Gaines and P. V. Facey, "Some experience in interactive system development and application." *Proc. IEEE*, **63**, 1975, pp. 894–911.
185. G. A. Miller and W. G. Madow, "On the maximum likelihood estimate of the Shannon-Wiener measure of information." AFCRC-TR-54-75, Operational Applications Laboratory, ARDC, August 1954.
186. W. R. Garner and W. J. McGill, "The relation between information and variance analyses." *Psychometrika*, **21**, pp. 219–228, 1956.
187. S. Kullback, *Information Theory and Statistics*. John Wiley, New York, 1959.
188. A. W. Biermann, "On the inference of Turing machines from sample computations." *Artificial Intelligence*, **3**, pp. 181–198, 1972.
189. A. W. Biermann, R. Baum, R. Krishnaswamy, and F. E. Petry, "Automatic program synthesis reports." OSU-CISRC-TR-73-6, Computer and Information Sciences Research Center, Ohio State University, USA, 1973.
190. S. Crespi-Reghizzi, M. A. Melankoff, and L. Lichten, "The use of grammatical inference for designing programming languages." *Comm. ACM*, **16**, 1973, pp. 83–90.
191. J. Hartmanis and R. E. Stearns, *Algebraic Structure Theory of Sequential Machines*. Prentice-Hall, New Jersey, 1966.
192. M. A. Arbib, *Algebraic Theory of Machines, Languages and Semigroups*. Academic Press, New York, 1968.
193. M. D. Mesarovic, *The Control of Multivariable Systems*. MIT Press, Mass., 1960.

194. M. D. Mesarovic, D. Macko, and Y. Takahara, *Theory of Hierarchical Multilevel Systems.* Academic Press, New York, 1970.
195. M. D. Mesarovic and Y. Takahara, *General Systems Theory: Mathematical Foundations.* Academic Press, New York, 1975.
196. G. J. Klir (ed.), *Trends in General Systems Theory.* John Wiley, New York, 1972.
197. B. P. Zeigler, "A conceptual basis for modelling and simulation." *Int. J. General Systems*, **1**, pp. 213–228, 1974.
198. B. P. Zeigler, "Simulation based structural complexity of models." *Int. J. General Systems*, **2**, pp. 217–223.
199. A. Koestler and J. R. Smythies (eds.), *Beyond Reductionism.* Hutchison, London, 1969.
200. R. Borger and F. Cioffi (eds.), *Explanation in the Behavioural Sciences.* Cambridge University Press, Cambridge, 1970.
201. J. Monod, *Chance and Necessity.* Collins, London, 1972.
202. J. Lewis (ed.), *Beyond Chance and Necessity.* Garnstone Press, London, 1974.
203. F. J. Ayala and T. Dobzhansky (eds.), *Studies in the Philosophy of Biology.* Macmillan, London, 1974.
204. T. W. Wann (ed.), *Behaviorism and Phenomenology.* University of Chicago Press, Chicago, 1964.
205. B. B. Wolman and E. Nagel (eds.), *Scientific Psychology.* Basic Books, New York, 1965.
206. N. S. Care and C. Landesman (eds.), *Readings in the Theory of Action.* Indiana University Press, Bloomington, 1968.
207. T. Mischel (ed.), *Human Action.* Academic Press, New York, 1969.
208. J. R. Royce (ed.), *Toward Unification in Psychology.* University of Toronto Press, Toronto, 1970.
209. S. C. Brown (ed.), *Philosophy of Psychology.* Macmillan, London, 1974.
210. A. Schutz, *The Phenomenology of the Social World* (trans. G. Walsh and F. Lehnert). Heinemann, 1972 (original 1932).
211. A. Schutz and T. Luckham, *The Structures of the Life-World* (trans. R. M. Zaner and H. T. Engelhardt). Heinemann, London, 1974.
212. H. M. Blalock (ed.), *Causal Models in the Social Sciences.* Macmillan, London, 1971.
213. S. Andreski (ed.), *The Essential Comte* (trans. M. Clarke), Croom Helm, London, 1974.
214. A. Giddens (ed.), *Positivism and Sociology.* Heinemann, London, 1974.
215. H. Feigl, M. Scriven and G. Maxwell (eds.), *Concepts, Theories and the Mind-Body Problem.* University of Minnesota Press, Minneapolis, 1958.
216. P. K. Feyerabend and G. Maxwell (eds.), *Mind, Matter and Method.* University of Minnesota Press, Minneapolis, 1966.
217. M. Radner and S. Winokur (eds.), *Analyses of the Theories and Methods of Physics and Psychology.* University of Minnesota Press, Mineapolis, 1970.
218. T. Shanin (ed.), *The Rules of the Game.* Tavistock Publications, London, 1972.
219. S. Korner (ed.), *Observation and Interpretation in the Philosophy of Physics.* Dover Publications, New York, 1957.
220. T. Bastin (ed.), *Quantum Theory and Beyond.* Cambridge University Press, Cambridge, 1971.
221. T. Gold (ed.), *The Nature of Time.* Cornell University Press, Ithaca, 1967.
222. R. M. Gale (ed.), *The Philosophy of Time.* Macmillan, 1968.
223. J. Zeman (ed.), *Time in Science and Philosophy.* Elsevier, Amsterdam, 1971.
224. N. P. Bhatia and G. P. Szego, *Dynamical Systems: Stability Theory and Applications.* Springer-Verlag, Berlin, 1967.
225. J. Auslander and W. H. Gottschalk (eds.), *Topological Dynamics.* W. A. Benjamin, New York, 1968.
226. L. Wittgenstein, *Tractatus Logico-Philosophicus* (trans. D. F. Pears and B. F. McGuiness), Routledge and Kegan Paul, London, 1972 (original 1921).
227. M. Merleau-Ponty, *The Primacy of Perception* (ed. J. M. Edie). Northwestern University Press, USA, 1964.
228. J. P. Sartre, *Being and Nothingness* (trans. H. E. Barnes), Methuen, London, 1958 (original 1943).
229. C. L. Becker, *The Heavenly City of the Eighteenth-Century Philosophers.* Yale University Press, New Haven, 1932.
230. J. Bradley, *Mach's Philosophy of Science.* Athlone Press, University of London, 1971.
231. E. Pivcevic, *Husserl and Phenomenology.* Hutchinson, London, 1970.
232. H. Spiegelberg, *The Phenomenological Movement, Vols. 1 and 2*, M. Nijhoff, The Hague, 1965.
233. M. Roche, *Phenomenology, Language and the Social Sciences.* Routledge and Kegan Paul, London, 1973.
234. L. Kolakowski, *The Alienation of Reason.* Doubleday, New York, 1968.
235. J. Ehrmann (ed.), *Structuralism.* Doubleday, New York, 1970.
236. D. Wilson, *Presuppositions and Non-Truth-Conditional Semantics.* Academic Press, London, 1975.
237. J. Giedymin, "Antipositivism in contemporary philosophy of social science and humanities." *Brit. J. Philos. Sci.*, **26**, 1975, pp. 275–301.
238. N. Koertge, "Popper's metaphysical research program for the human sciences." *Inquiry*, **18**, 1975, pp. 437–462.
239. E. Erwin, "The confirmation machine." In Ref. 11, pp. 306–321.
240. G. J. Klir, "Computerized methodology for structure modelling." *Annals of Systems Research*, Vol. 4. H. E. Stenfert Kroese, Leiden, the Netherlands, 1976.

For Prof. Gaines' biography, see Vol. 2, No. 4, p. 208 of this journal.

Downloaded by [University of Calgary] at 10:52 14 February 2012