Reprinted from JOURNAL OF MATHEMATICAL PSYCHOLOGY Vol. 6, No. 3, October 1969
Printed in Belgium

Stimulus–Response Theory of Finite Automata[1]

PATRICK SUPPES

Stanford University, Stanford, California 94305

The central aim of the paper is to state and prove a representation theorem for finite automata in terms of models of stimulus–response theory. The main theorem is that, given any connected finite automaton, there is a stimulus–response model that asymptotically becomes isomorphic to it. Implications of this result for language learning are discussed in some detail. In addition, an immediate corollary is that any tote hierarchy in the sense of Miller and Chomsky is isomorphic to some stimulus–response model at asymptote. Representations of probabilistic automata are also discussed, and an application to the learning of arithmetic algorithms is given.

1. INTRODUCTION

Ever since the appearance of Chomsky's famous review (1959) of Skinner's *Verbal Behavior* (1957), linguists have conducted an effective and active campaign against the empirical or conceptual adequacy of any learning theory whose basic concepts are those of stimulus and response, and whose basic processes are stimulus conditioning and stimulus sampling.

Because variants of stimulus–response theory had dominated much of experimental psychology in the two decades prior to the middle fifties, there is no doubt that the attack of the linguists has had a salutary effect in disturbing the theoretical complacency of many psychologists. Indeed, it has posed for all psychologists interested in systematic theory a number of difficult and embarrassing questions about language learning and language behavior in general. However, in the flush of their initial victories, many linguists have made extravagant claims and drawn sweeping, but unsupported conclusions about the inadequacy of stimulus–response theories to handle any central aspects of language behavior. I say "extravagant" and "unsupported" for this reason. The claims and conclusions are supported neither by careful mathematical argument to show that in principle a conceptual inadequacy is to be found in all standard stimulus–response theories, nor by systematic presentation of empirical evidence to show that the basic assumptions of these theories are empirically false. To cite two recent books of some importance, neither theorems nor data are to be found in

[1] This research has been partially supported by the National Science Foundation and the U. S. Office of Education. I am indebted to Kenneth N. Wexler for a number of helpful comments on the manuscript.

Chomsky (1965) or Katz and Postal (1964), but rather one can find many useful examples of linguistic analysis, many interesting and insightful remarks about language behavior, and many incompletely worked out arguments about theories of language learning.

The central aim of the present paper and its projected successors is to prove in detail that stimulus–response theory, or at least a mathematically precise version, can indeed give an account of the learning of many phrase-structure grammars. I hope that there will be no misunderstanding about the claims I am making. The mathematical definitions and theorems given here are entirely subservient to the conceptual task of showing that the basic ideas of stimulus–response theory are rich enough to generate in a natural way the learning of many phrase-structure grammars. I am not claiming that the mathematical constructions in this paper correspond in any exact way to children's actual learning of their first language or to the learning of a second language at a later stage. A number of fundamental empirical questions are generated by the formal developments in this paper, but none of the relevant investigations have yet been carried out. Some suggestions for experiments are mentioned below. I have been concerned to show that linguists are quite mistaken in their claims that even *in principle*, apart from any questions of empirical evidence, it is not possible for conditioning theory to give an account of any essential parts of language learning. The main results in this paper, and its sequel dealing with general context-free languages, show that this linguistic claim is false. The specific constructions given below show that linguistic objections to the processes of stimulus-conditioning and sampling as being unable in principle to explain any central aspects of learning a grammar must be reformulated in less sweeping generality.

The mathematical formulation and proof of the main results presented here require the development of a certain amount of formal machinery. In order not to obscure the main ideas, it seems desirable to describe in a preliminary and intuitive fashion the character of the results.

The central idea is quite simple—it is to show how by applying accepted principles of conditioning an organism may theoretically be taught by an appropriate reinforcement schedule to respond as a finite automaton. An automaton is defined as a device with a finite number of internal states. When it is presented with one of a finite number of letters from an alphabet, as a function of this letter of the alphabet and its current internal state, it moves to another one of its internal states. (A more precise mathematical formulation is given below.) In order to show that an organism obeying general laws of stimulus conditioning and sampling can be conditioned to become an automaton, it is necessary first of all to interpret within the usual run of psychological concepts, the notion of a letter of an alphabet and the notion of an internal state. In my own thinking about these matters, I was first misled by the perhaps natural attempt to identify the internal state of the automaton with the state of conditioning of the organism. This idea, however, turned out to be clearly wrong.

In the first place, the various possible states of conditioning of the organism correspond to various possible automata that the organism can be conditioned to become. Roughly speaking, to each state of conditioning there corresponds a different automaton. Probably the next most natural idea is to look at a given conditioning state and use the conditioning of individual stimuli to represent the internal states of the automaton. In very restricted cases this correspondence works, but in general it does not, for reasons that become clear below. The correspondence that turns out to work is the following: the internal states of the automaton are identified with the responses of the organism. There is no doubt that this "surface" behavioral identification will make many linguists concerned with deep structures (and other deep, abstract ideas) uneasy, but fortunately it is an identification already suggested in the literature of automata theory by E. F. Moore and others. The suggestion was originally made to simplify the formal characterization of automata by postulating a one–one relation between internal states of the machine and outputs of the machine. From a formal standpoint this means that the two separate concepts of internal state and output can be welded into the single concept of internal state and, for our purposes, the internal states can be identified with responses of the organism.

The correspondence to be made between letters of the alphabet that the automaton will accept and the appropriate objects within stimulus–response theory is fairly obvious. The letters of the alphabet correspond in a natural way to sets of stimulus elements presented on a given trial to an organism. So again, rougly speaking, the correspondence in this case is between the alphabet and selected stimuli. It may seem like a happy accident, but the correspondences between inputs to the automata and stimuli presented to the organism, and between internal states of the machine and responses of the organsim, are conceptually very natural.

Because of the conceptual importance of the issues that have been raised by linguists for the future development of psychological theory, perhaps above all because language behavior is the most characteristically human aspect of our behavior patterns, it is important to be as clear as possible about the claims that can be made for a stimulus–response theory whose basic concepts seem so simple and to many so woefully inadequate to explain complex behavior, including language behavior. I cannot refrain from mentioning two examples that present very useful analogies. First is the reduction of all standard mathematics to the concept of set and the simple relation of an element being a member of a set. From a naive standpoint, it seems unbelievable that the complexities of higher mathematics can be reduced to a relation as simple as that of set membership. But this is indubitably the case, and we know in detail how the reduction can be made. This is not to suggest, for instance, that in thinking about a mathematical problem or even in formulating and verifying it explicitly, a mathematician operates simply in terms of endlessly complicated statements about set membership. By appropriate explicit definition we introduce many additional concepts, the ones actually used in discourse. The fact remains, however, that the reduction to the single relationship of set membership can be made and in fact has been carried out in detail. The second example, which is close to our present inquiry, is the status of simple machine languages for computers. Again, from the naive standpoint it seems incredible that modern computers can do the things they can in terms either of information processing or numerical computing when their basic language consists essentially just of finite sequences of 1's and 0's; but the more complex computer languages that have been introduced are not at all for the convenience of the machines but for the convenience of human users. It is perfectly clear how any more complex language, like ALGOL, can be reduced by a compiler or other device to a simple machine language. The same attitude, it seems to me, is appropriate toward stimulus–response theory. We cannot hope to deal directly in stimulus–response connections with complex human behavior. We can hope, as in the two cases just mentioned, to construct a satisfactory systematic theory in terms of which a chain of explicit definitions of new and ever more complex concepts can be introduced. It is these new and explicitly defined concepts that will be related directly to the more complex forms of behavior. The basic idea of stimulus–response association or connection is close enough in character to the concept of set membership or to the basic idea of automata to make me confident that new and better versions of stimulus–response theory may be expected in the future and that the scientific potentiality of theories stated essentially in this framework has by no means been exhausted.

Before turning to specific mathematical developments, it will be useful to make explicit how the developments in this paper may be used to show that many of the common conceptions of conditioning, and particularly the claims that conditioning refers only to simple reflexes like those of salivation or eye blinking, are mistaken. The mistake is to confuse particular restricted applications of the fundamental theory with the range of the theory itself. Experiments on classical conditioning do indeed represent a narrow range of experiments from a broader conceptual standpoint. It is important to realize, however, that *experiments* on classical conditioning do not define the range and limits of conditioning theory itself. The main aim of the present paper is to show how any finite automaton, no matter how complicated, may be constructed purely within stimulus–response theory. But from the standpoint of automata, classical conditioning represents a particularly trivial example of an automaton. Classical conditioning may be represented by an automaton having a one-letter alphabet and a single internal state. The next simplest case corresponds to the structure of classical discrimination experiments. Here there is more than a single letter to the alphabet, but the transition table of the automaton depends in no way on the internal state of the automaton. In the case of discrimination, we may again think of the responses as corresponding to the internal states of the automaton. In this sense there is more than one internal state, contrary to the case of classical conditioning, but what is fundamental is that the transition table of the automaton does not depend on the internal states but only on the external stimuli presented according to a schedule fixed by the experimenter. It is of the utmost importance to realize that this restriction, as in the case of classical conditioning experiments, is not a restriction that is in any sense inherent in conditioning theory itself. It merely represents concentration on a certain restricted class of experiments.

Leaving the technical details for later, it is still possible to give a very clear example of conditioning that goes beyond the classical cases and yet represents perhaps the simplest nontrivial automaton. By nontrivial I mean: there is more than one letter in the alphabet; there is more than one internal state; and the transition table of the automaton is a function of both the external stimulus and the current internal state. As an example, we may take a rat being run in a maze. The reinforcement schedule for the rat is set up so as to make the rat become a two-state automaton. We will use as the external alphabet of the automaton a two-letter alphabet consisting of a black or a white card. Each choice point of the maze will consist of either a left turn or a right turn. At each choice point either a black card or a white card will be present. The following table describes both the reinforcement schedule and the transition table of the automaton.

	L	R
LB	1	0
LW	0	1
RB	0	1
RW	1	0

Thus the first row shows that when the previous response has been left (L) and a black stimulus card (B) is presented at the choice point, with probability one the animal is reinforced to turn left. The second row indicates that when the previous response is left and a white stimulus card is presented at the choice point, the animal is reinforced 100% of the time to turn right, and so forth, for the other two possibilities. From a formal standpoint this is a simple schedule of reinforcement, but already the double aspect of contingency on both the previous response and the displayed stimulus card makes the schedule more complicated in many respects than the schedules of reinforcement that are usually run with rats. I have not been able to get a uniform prediction from my experimental colleagues as to whether it will be possible to teach rats to learn this schedule. (Most of them are confident pigeons can be trained to respond like nontrivial two-state automata.)[2] One thing to note about this schedule is that it is recursive in the sense that if the animal is properly trained according to the schedule, the length of the maze will be of no importance. He will always make a

[2] Indeed, Phoebe C. E. Diebold and Ebbey Bruce Ebbesen have already successfully trained two pigeons. Sequences of several hundred responses are easily obtained, and the error rate is surprisingly low—well under 1%.

response that depends only upon his previous response and the stimulus card present at the choice point.

There is no pretense that this simple two-state automaton is in any sense adequate to serious learning. I am not proposing, for example, that there is much chance of teaching even a simple one-sided linear grammar to rats. I am proposing to psychologists, however, that already automata of a small number of states present immediate experimental challenges in terms of what can be done with animals of each species. For example, what is the most complicated automaton a monkey may be trained to imitate? In this case, there seems some possibility of approaching at least reasonably complex one-sided linear grammars (using the theorem that any one-sided linear grammar is definable by a finite-state automaton). In the case of the lower species, it will be necessary to exploit to the fullest the kind of stimuli to which the organisms are most sensitive and responsive in order to maximize the complexity of the automata they can imitate.

If a generally agreed upon definition of complexity for finite automata can be reached, it will be possible to use this measure to gauge the relative level of organizational complexity that can be achieved by a given species, at least in terms of an external schedule of conditioning and reinforcement. I do want to emphasize that the measures appropriate to experiments with animals are almost totally different from the measures that have been discussed in the recent literature of automata as complexity measures for computations. What is needed in the case of animals is the simple and orderly arrangement on a complexity scale of automata that have a relatively small number of states and that accept a relatively small alphabet of stimuli. The number of distinct training conditions is not a bad measure and can be used as a first approximation. Thus in the case of classical conditioning, this number is one. In the case of discrimination between black and white stimuli, the number is two. In the case of the two-state automaton described for the maze experiment, this number is four, but there are some problems with this measure. It is not clear that we would regard as more complex than this two-state automaton, an organism that masters a discrimination experiment consisting of six different responses to six different discriminating stimuli. Consequently, what I've said here about complexity is pre-systematic. I do think the development of an appropriate scale of complexity can be of theoretical interest, especially in cross-species comparison of intellectual power.

The remainder of this paper is devoted to the technical development of the general ideas already discussed. Section 2 is concerned with standard notions of finite and probabilistic automata. Readers already familiar with this literature should skip this section and go on to the treatment of stimulus–response theory in Sec. 3. It has been necessary to give a rigorous axiomatization of stimulus–response theory in order to formulate the representation theorem for finite automata in mathematically precise form. However, the underlying ideas of stimulus–response theory as formulated in Sec. 3 will be familiar to all experimental psychologists. In Sec. 4 the most important result of the paper is proved, namely, that any finite automaton can be represented at asymptote by an appropriate model of stimulus–response theory. In Sec. 5 some extensions of these results to probabilistic automata are sketched, and an example from arithmetic is worked out in detail.

The relationship between stimulus–response theory and grammars is established in Sec. 4 by known theorems relating automata to grammars. The results in the present paper are certainly restricted regarding the full generality of context-free languages. Weakening these restrictions will be the focus of a subsequent paper.

Some results on tote hierarchies and *plans* in the sense of Miller, Galanter, and Pribram (1960) are also given in Sec. 4. The representation of tote hierarchies by stimulus–response models follows directly from the main theorem of that section.

2. AUTOMATA

The account of automata given here is formally self-contained, but not really self-explanatory in the deeper sense of discussing and interpreting in adequate detail the systematic definitions and theorems. I have followed closely the development in the well-known article of Rabin and Scott (1959), and for probabilistic automata, the article of Rabin (1963).

DEFINITION 1. *A structure* $\mathfrak{A} = \langle A, \Sigma, M, s_0, F\rangle$ *is a finite (deterministic) automaton if and only if*

(i) *A is a finite, nonempty set (the set of states of* $\mathfrak{A}$),

(ii) *Σ is a finite, nonempty set (the alphabet)*,

(iii) *M is a function from the Cartesian product $A \times \Sigma$ to A (M defines the transition table of* $\mathfrak{A}$),

(iv) s_0 *is in A (s_0 is the initial state of* $\mathfrak{A}$),

(v) *F is a subset of A (F is the set of final states of* $\mathfrak{A}$).

In view of the generality of this definition it is apparent that there are a great variety of automata, but as we shall see, this generality is easily matched by the generality of the models of stimulus–response theory.

In notation now nearly standardized, Σ^* is the set of finite sequences of elements of Σ, including the empty sequence Λ. The elements of Σ^* are ordinarily called *tapes*. If $\sigma_1, \ldots, \sigma_k$ are in Σ, then $x = \sigma_1 \cdots \sigma_k$ is in Σ^*. (As shown below, these tapes correspond in a natural way to finite sequences of sets of stimulus elements.) The function M can be extended to a function from $A \times \Sigma^*$ to A by the following recursive definition for s in A, x in Σ^*, and σ in Σ.

$$M(s, \Lambda) = s,$$

$$M(s, x\sigma) = M(M(s, x), \sigma).$$

Definition 2. *A tape x of Σ^* is accepted by $\mathfrak{A}$ if and only if $M(s_0, x)$ is in F. A tape x that is accepted by $\mathfrak{A}$ is a sentence of $\mathfrak{A}$.*

We shall also refer to tapes as *strings* of the alphabet Σ.

Definition 3. *The language L generated by $\mathfrak{A}$ is the set of all sentences of $\mathfrak{A}$, i.e., the set of all tapes accepted by $\mathfrak{A}$.*

Regular languages are sometimes defined just as those languages generated by some finite automaton. An independent, set-theoretical characterization is also possible. The basic result follows from Kleene's (1956) fundamental analysis of the kind of events definable by McCulloch-Pitts nets. Several equivalent formulations are given in the article by Rabin and Scott (1959). From a linguistic standpoint probably the most useful characterization is that to be found in Chomsky (1963, pp. 368–371). Regular languages are generated by one-sided linear grammars. Such grammars have a finite number of rewrite rules, which in the case of right-linear rules, are of the form

$$A \to xB.$$

Whichever of several equivalent formulations is used, the fundamental theorem, originally due to Kleene, but closely related to the theorem of Myhill given by Rabin and Scott, in this.

Theorem on Regular Languages. *Any regular language is generated by some finite automaton, and every finite automaton generates a regular language.*

For the main theorem of this article, we need the concepts of isomorphism and equivalence of finite automata. The definition of isomorphism is just the natural set-theoretical one for structures like automata.

Definition 4. *Let $\mathfrak{A} = \langle A, \Sigma, M, s_0, F\rangle$ and $\mathfrak{A}' = \langle A', \Sigma', M', s_0', F'\rangle$ be finite automata. Then $\mathfrak{A}$ and $\mathfrak{A}'$ are isomorphic if and only if there exists a function f such that*

(i) *f is one–one,*

(ii) *Domain of f is $A \cup \Sigma$ and range of f is $A' \cup \Sigma'$,*

(iii) *For every a in $A \cup \Sigma$*

$$a \in A \qquad \text{if and only if} \qquad f(a) \in A',$$

(iv) *For every s in A and σ in Σ*

$$f(M(s, \sigma)) = M'(f(s), f(\sigma)),$$

(v) $f(s_0) = s_0'$,

(vi) *For every s in A*

$$s \in F \qquad \text{if and only if} \qquad f(s) \in F'.$$

It is apparent that conditions (i)–(iii) of the definition imply that for every a in $A \cup \Sigma$

$$a \in \Sigma \qquad \text{if and only if} \qquad f(a) \in \Sigma',$$

and consequently, this condition on Σ need not be stated. From the standpoint of the general algebraic or set-theoretical concept of isomorphism, it would have been more natural to define an automaton in terms of a basic set $B = A \cup \Sigma$, and then require that A and Σ both be subsets of B. Rabin and Scott avoid the problem by not making Σ a part of the automaton. They define the concept of an automaton $\mathfrak{A} = \langle A, M, s_0, F\rangle$ with respect to an alphabet Σ, but for the purposes of this paper it is also desirable to include the alphabet Σ in the definition of $\mathfrak{A}$ in order to make explicit the natural place of the alphabet in the stimulus–response models, and above all, to provide a simple setup for going from one alphabet Σ to another Σ'. In any case, exactly how these matters are handled is not of central importance here.

DEFINITION 5. *Two automata are equivalent if and only if they accept exactly the same set of tapes.*

This is the standard definition of equivalence in the literature. As it stands, it means that the definition of equivalence is neither stronger nor weaker than the definition of isomorphism, because, on the one hand, equivalent automata are clearly not necessarily isomorphic, and, on the other hand, isomorphic automata with different alphabets are not equivalent. It would seem natural to weaken the notion of equivalence to include two automata that generate distinct but isomorphic languages, or sets of tapes, but this point will bear on matters here only tangentially.

A finite automaton is *connected* if for every state s there is a tape x such that $M(s_0, x) = s$. It is easy to show that every automaton is equivalent to a connected automaton, and the representation theorem of Sec. 4 is restricted to connected automata. It is apparent that from a functional standpoint, states that cannot be reached by any tape are of no interest, and consequently, restriction to connected automata does not represent any real loss of generality. The difficulty of representing automata with unconnected states by stimulus–response models is that we have no way to condition the organism with respect to these states, at least in terms of the approach developed here.

It is also straightforward to establish a representation for probabilistic automata within stimulus–response theory, and, as is apparent in Sec. 5, there are some interesting differences in the way we may represent deterministic and probabilistic automata within stimulus–response theory.

DEFINITION 6. *A structure $\mathfrak{A} = \langle A, \Sigma, p, s_0, F\rangle$ is a (finite) probabilistic automaton if and only if*

(i) *A is a finite, nonempty set,*

(ii) *Σ is a finite, nonempty set,*

(iii) *p is a function on $A \times \Sigma$ such that for each s in A and σ in Σ, $p_{s,\sigma}$ is a probability density over A*, i.e.,

(a) *for each s' in A, $p_{s,\sigma}(s') \geqslant 0$,*

(b) $\sum_{s' \in A} p_{s,\sigma}(s') = 1$,

(iv) *s_0 is in A,*

(v) *F is a subset of A.*

The only change in generalizing from Definition 1 to Definition 6 is found in (iii), although it is natural to replace (iv) by an initial probability density. It is apparent how Definition 4 must be modified to characterize the isomorphism of probabilistic automata, and so the explicit definition will not be given.

3. STIMULUS-RESPONSE THEORY

The formalization of stimulus–response theory given here follows closely the treatment in Estes and Suppes (1959) and Suppes and Atkinson (1960). Some minor changes have been made to facilitate the treatment of finite automata, but it is to be strongly emphasized that none of the basic ideas or assumptions has required modification.

The theory is based on six primitive concepts, each of which has a direct psychological interpretation. The first one is the set S of stimuli, which we shall assume is not empty, but which we will not restrict to being either finite or infinite on all occasions. The second primitive concept is the set R of responses and the third primitive concept the set E of possible reinforcements. As in the case of the set of stimuli, we need not assume that either R or E is finite, but in the present applications to the theory of finite automata we shall make this restrictive assumption. (For a proper treatment of phonology it will clearly be necessary to make R, and probably E as well, infinite with at the very least a strong topological if not metric structure.)

The fourth primitive concept is that of a measure μ on the set of stimuli. In case the set S is finite this measure is often the number of elements in S. For the general theory we shall assume that the measure of S itself is always finite, i.e., $\mu(S) < \infty$.

The fifth primitive concept is the sample space X. Each element x of the sample space represents a possible experiment, that is, an infinite sequence of trials. In the present theory, each trial may be described by an ordered quintuple $\langle C, T, s, r, e \rangle$, where C is the conditioning function, T is the subset of stimuli presented to the organism on the given trial, s is the sampled subset of T, r is the response made on the trial, and e is the reinforcement occurring on that trial. It is not possible to make all the comments here that are required for a full interpretation and understanding of the theory. For those wanting a more detailed description, the two references already given will prove useful. A very comprehensive set of papers on stimulus-sampling theory has been put together in the collection edited by Neimark and Estes (1967). The present version of stimulus–response theory should in many respects be called stimulus-sampling theory, but I have held to the more general stimulus–response terminology to emphasize the juxtaposition of the general ideas of behavioral psychology on the one hand and linguistic theory on the other. In addition, in the theoretical applications to be made here the specific sampling aspects of stimulus–response theory are not as central as in the analysis of experimental data.

Because of the importance to be attached later to the set T of stimuli presented on each trial, its interpretation in classical learning theory should be explicitly mentioned. In the case of simple learning, for example, in classical conditioning, the set T is the same on all trials and we would ordinarily identify the sets T and S. In the case of discrimination learning, the set T varies from trial to trial, and the application we are making to automata theory falls generally under the discrimination case. The conditioning function C is defined over the set R of responses and C_r is the subset of S conditioned or connected to response r on the given trial. How the conditioning function changes from trial to trial is made clear by the axioms.

From the quintuple description of a given trial it is clear that certain assumptions about the behavior that occurs on a trial have already been made. In particular it is apparent that we are assuming that only one sample of stimuli is drawn on a given trial, that exactly one response occurs on a trial and that exactly one reinforcement occurs on a trial. These assumptions have been built into the set-theoretical description of the sample space X and will not be an explicit part of our axioms.

Lying behind the formality of the ordered quintuples representing each trial is the intuitively conceived temporal ordering of events on any trial, which may be represented by the following diagram:

State of conditioning at beginning of trial	→	presentation of stimuli	→	sampling of stimuli	→	response	→	reinforcement	→	state of conditioning at beginning of new trial
C_n	→	T_n	→	s_n	→	r_n	→	e_n	→	C_{n+1}

The sixth and final primitive concept is the probability measure P on the appropriate Borel field of cylinder sets of X. The exact description of this Borel field is rather complicated when the set of stimuli is not finite, but the construction is standard, and we shall assume the reader can fill in details familiar from general probability theory. It is emphasized that all probabilities must be defined in terms of the measure P.

We also need certain notation to take us back and forth between elements or subsets of the sets of stimuli, responses, and reinforcements to events of the sample space X.

First, r_n is the event of response r on trial n, that is, the set of all possible experimental realizations or elements of X having r as a response on the nth trial. Similarly, $e_{r,n}$ is the event of response r's being reinforced on trial n. The event $e_{0,n}$ is the event of no reinforcement on trial n. In like fashion, C_n is the event of conditioning function C occurring on trial n, T_n is the event of presentation set T occurring on trial n, and so forth. Additional notation that does not follow these conventions will be explicitly noted.

We also need a notation for sets defined by events occurring up to a given trial. Reference to such sets is required in expressing that central aspects of stimulus conditioning and sampling are independent of the pattern of past events. If I say that Y_n is an n-cylinder set, I mean that the definition of Y_n does not depend on any event occurring after trial n. However, an even finer breakdown is required that takes account of the postulated sequence $C_n \rightarrow T_n \rightarrow s_n \rightarrow r_n \rightarrow e_n$ on a given trial, so in saying that Y_n is a C_n-cylinder set what is meant is that its definition does not depend on any event occurring after C_n on trial n, i.e., its definition could depend on T_{n-1} or C_n, for example, but not on T_n or s_n. As an abbreviated notation, I shall write $Y(C_n)$ for this set and similarly for other cylinder sets. The notation Y_n without additional qualification shall always refer to an n-cylinder set.

Also, to avoid an overly cumbersome notation, event notation of the sort already indicated will be used, e.g., $e_{r,n}$, for reinforcement of response r on trial n, but also the notation $\sigma \in C_{r,n}$ for the event of stimulus σ's being conditioned to response r on trial n.

To simplify the formal statement of the axioms it is assumed without repeated explicit statement that any given events on which probabilities are conditioned have positive probability. Thus, for example, the tacit hypothesis of Axiom S2 is that $P(T_m) > 0$ and $P(T_n) > 0$.

The axioms naturally fall into three classes. Stimuli must be sampled in order to be conditioned, and they must be conditioned in order for systematic response patterns to develop. Thus, there are naturally three kinds of axioms: sampling axioms; conditioning axioms; and response axioms. A verbal formulation of each axiom is given together with its formal statement. From the standpoint of formulations of the theory already in the literature, perhaps the most unusual feature of the present axioms is not to require that the set S of stimuli be finite. It should also be emphasized that for any one specific kind of detailed application additional specializing assumptions are needed. Some indication of these will be given in the particular application to automata theory, but it would take us too far afield to explore these specializing assumptions in any detail and with any faithfulness to the range of assumptions needed for different experimental applications.

DEFINITION 7. *A structure* $\mathscr{S} = \langle S, R, E, \mu, X, P \rangle$ *is a stimulus–response model if and only if the following axioms are satisfied:*

SAMPLING AXIOMS.

S1. $P(\mu(s_n) > 0) = 1.$

(*On every trial a set of stimuli of positive measure is sampled with probability* 1.)

S2. $P(s_m \mid T_m) = P(s_n \mid T_n).$

(*If the same presentation set occurs on two different trials, then the probability of a given sample is independent of the trial number.*)

S3. *If* $s \cup s' \subseteq T$ *and* $\mu(s) = \mu(s')$ *then* $P(s_n \mid T_n) = P(s_n' \mid T_n).$

(*Samples of equal measure that are subsets of the presentation set have an equal probability of being sampled on a given trial.*)

S4. $P(s_n \mid T_n , Y_n(C_n)) = P(s_n \mid T_n).$

(*The probability of a particular sample on trial n, given the presentation set of stimuli, is independent of any preceding pattern* $Y_n(C_n)$ *of events.*)

CONDITIONING AXIOMS:

C1. *If* $r, r' \in R$, $r \neq r'$ *and* $C_r \cap C_{r'} \neq 0$, *then* $P(C_n) = 0.$

(*On every trial with probability* 1 *each stimulus element is conditioned to at most one response.*)

C2. *There exists a* $c > 0$ *such that for every* σ, C, r, n, s, e_r, *and* Y_n

$$P(\sigma \in C_{r,n+1} \mid \sigma \notin C_{r,n} , \sigma \in s_n , e_{r,n} , Y_n) = c.$$

(*The probability is c of any sampled stimulus element's becoming conditioned to the reinforced response if it is not already so conditioned, and this probability is independent of the particular response, trial number, or any preceding pattern* Y_n *of events.*)

C3. $P(C_{n+1} \mid C_n , e_{0,n}) = 1.$

(*With probability* 1, *the conditioning of all stimulus elements remains the same if no response is reinforced.*)

C4. $P(\sigma \in C_{r,n+1} \mid \sigma \in C_{r,n} , \sigma \notin s_n , Y_n) = 1.$

(*With probability* 1, *the conditioning of unsampled stimuli does not change.*)

RESPONSE AXIOMS:

R1. *If* $\bigcup_{r \in R} C_r \cap s \neq 0$ *then*

$$P(r_n \mid C_n , s_n , Y(s_n)) = \frac{\mu(s \cap C_r)}{\mu(s \cap UC_r)} .$$

(*If at least one sampled stimulus is conditioned to some response, then the probability of any response is the ratio of the measure of sampled stimuli conditioned to this response* *to the measure of all the sampled conditioned stimuli, and this probability is independent of any preceding pattern* $Y(s_n)$ *of events.*

R2. *If* $\bigcup_{r \in R} C_r \cap s = 0$ *then there is a number* ρ_r *such that*

$$P(r_n \mid C_n, s_n, Y(s_n)) = \rho_r .$$

(*If no sampled stimulus is conditioned to any response, then the probability of any response* r *is a constant guessing probability* ρ_r *that is independent of* n *and any preceding pattern* $Y(s_n)$ *of events.*)

A general discussion of these axioms and their implications for a wide range of psychological experiments may be found in the references already cited. The techniques of analysis used in Sec. 5 of the present paper are extensively exploited and applied to a number of experiments in Suppes and Atkinson (1960).

4. REPRESENTATION OF FINITE AUTOMATA

A useful beginning for the analysis of how we may represent finite automata by stimulus–response models is to examine what is wrong with the most direct approach possible. The difficulties that turn up may be illustrated by the simple example of a two-letter alphabet (i.e., two stimuli σ_1 and σ_2, as well as the "start-up" stimulus σ_0) and a two-state automaton (i.e., two responses r_1 and r_2). Consideration of this example, already mentioned in the introductory section, will be useful for several points of later discussion.

By virtue of Axiom S1, the single presented stimulus must be sampled on each trial, and we assume that for every n,

$$0 < P(\sigma_0 \in s_n),\ P(\sigma_1 \in s_n),\ P(\sigma_2 \in s_n) < 1.$$

Suppose, further, the transition table of the machine is:

	r_1	r_2
$r_1\sigma_1$	1	0
$r_1\sigma_2$	0	1
$r_2\sigma_1$	0	1
$r_2\sigma_2$	1	0

which requires knowledge of both r_i and σ_j to predict what response should be next. The natural and obvious reinforcement schedule for imitating this machine is:

$$P(e_{1,n} \mid \sigma_{1,n}, r_{1,n-1}) = 1,$$
$$P(e_{2,n} \mid \sigma_{1,n}, r_{2,n-1}) = 1,$$

$$P(e_{2,n} \mid \sigma_{2,n}, r_{1,n-1}) = 1,$$
$$P(e_{1,n} \mid \sigma_{2,n}, r_{2,n-1}) = 1,$$

where $\sigma_{i,n}$ is the event of stimulus σ_i's being sampled on trial n. But for this reinforcement schedule the conditioning of each of the two stimuli continues to fluctuate from trial to trial, as may be illustrated by the following sequence. For simplification and without loss of generality, we may assume that the conditioning parameter c is 1, and we need indicate no sampling, because as already mentioned, the single stimulus element in each presentation set will be sampled with probability 1. We may represent the states of conditioning (granted that each stimulus is conditioned to either r_1 or r_2) by subsets of $S = \{\sigma_0, \sigma_1, \sigma_2\}$. Thus, if $\{\sigma_1, \sigma_2\}$ represents the conditioning function, this means both elements σ_1 and σ_2 are conditioned to r_1; $\{\sigma_1\}$ means that only σ_1 is conditioned to r_1, and so forth. Consider then the following sequence from trial n to $n + 2$:

$$\langle\{\sigma_2\}, \sigma_2 \in s, r_1, e_2\rangle \rightarrow \langle 0, \sigma_2 \in s, r_2, e_2\rangle \rightarrow \langle 0, \sigma_2 \in s, r_2, e_1\rangle.$$

The response on trial $n + 1$ satisfies the machine table, but already on $n + 2$ it does not, for $r_{2,n-1}\sigma_{2,n-2}$ should be followed by $r_{1,n-2}$. It is easy to show that this difficulty is fundamental and arises for any of the four possible conditioning states. (In working out these difficulties explicitly, the reader should assume that each stimulus is conditioned to either r_1 or r_2, which will be true for n much larger than 1 and $c = 1$.)

What is needed is a quite different definition of the states of the Markov chain of the stimulus–response model. [For proof of a general Markov-chain theorem for stimulus–response theory, see Estes and Suppes (1959).] Naively, it is natural to take as the states of the Markov chain the possible states of conditioning of the stimuli in S, but this is wrong on two counts in the present situation. First, we must condition the *patterns* of responses and presentation sets, so we take as the set of stimuli for the model, $R \times S$, i.e., the Cartesian product of the set R of responses and the set S of stimuli. What the organism must be conditioned to respond to on trial n is the pattern consisting of the preceding response given on trial $n - 1$ and the presentation set occurring on trial n.

It is still not sufficient to define the states of the Markov chain in terms of the states of conditioning of the elements in $R \times S$, because for reasons that are given explicitly and illustrated by many examples in Estes and Suppes (1959) and Suppes and Atkinson (1960), it is also necessary to include in the definition of state the response r_i that actually occurred on the preceding trial. The difficulty that arises if $r_{i,n-1}$ is not included in the definition of state may be brought out by attempting to draw the tree in the case of the two-state automaton already considered. Suppose just the pattern $r_1\sigma_1$ is conditioned, and the other four patterns, σ_0, $r_1\sigma_2$, $r_2\sigma_1$, and $r_2\sigma_2$, are not. Let us represent this conditioning state by C_1, and let τ_j be the noncontingent probability of σ_j, $0 \leqslant j \leqslant 2$, on every trial with every $\tau_j > 0$. Then the tree looks like this.

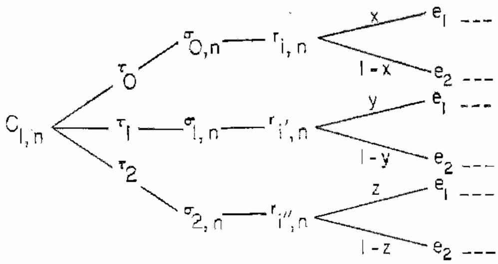

The tree is incomplete, because without knowing what response actually occurred on trial $n - 1$ we cannot complete the branches (e.g., specify the responses), and for a similar reason we cannot determine the probabilities x, y, and z. Moreover, we cannot remedy the situation by including among the branches the possible responses on trial $n - 1$, for to determine their probabilities we would need to look at trial $n - 2$, and this regression would not terminate until we reached trial 1.

So we include in the definition of state the response on trial $n - 1$. On the other hand, it is not necessary in the case of deterministic finite automata to permit among the states all possible conditioning of the patterns in $R \times S$. We shall permit only two possibilities—the pattern is unconditioned or it is conditioned to the appropriate response because conditioning to the wrong response occurs with probability zero. Thus with p internal states or responses and m letters in Σ, there are $(m + 1)\,p$ patterns, each of which is in one of two states, conditioned or unconditioned, and there are p possible preceding responses, so the number of states in the Markov chain is $p2^{(m+1)p}$. Actually, it is convenient to reduce this number further by treating σ_0 as a single pattern regardless of what preceding response it is paired with. The number of states is then $p2^{mp+1}$. Thus, for the simplest 2-state, 2-alphabet automaton, the number of states is 64. We may denote the states by ordered $mp + 2$-tuples

$$\langle r_j , i_0 , i_{0,1} ,..., i_{0,m} ,..., i_{p-1,m} \rangle,$$

where i_{kl} is 0 or 1 depending on whether the pattern $r_k\sigma_l$ is unconditioned or conditioned with $0 \leqslant k \leqslant p - 1$ and $1 \leqslant \ell \leqslant m$; r_j is the response on the preceding trial, and i_0 is the state of conditioning of σ_0. What we want to prove is that starting in the purely unconditioned set of states $\langle r_j , 0, 0,..., 0\rangle$, with probability 1 the system will always ultimately be in a state that is a member of the set of fully conditioned states $\langle r_j , 1, 1,..., 1\rangle$. The proof of this is the main part of the proof of the basic representation theorem.

Before turning to the theorem we need to define explicitly the concept of a stimulus–response model's asymptotically becoming an automaton. As has already been suggested, an important feature of this definition is this. The basic set S of stimuli corresponding to the alphabet Σ of the automaton is not the basic set of stimuli of the stimulus–response model, but rather, this basic set is the Cartesian product $R \times S$, where R is the set of responses. Moreover, the definition has been framed in such a way as to permit only a single element of S to be presented and sampled on each trial; this, however, is an inessential restriction used here in the interest of conceptual and notational simplicity. Without this restriction the basic set would be not $R \times S$, but $R \times P(S)$, where $P(S)$ is the power set of S, i.e., the set of all subsets of S, and then each letter of the alphabet Σ would be a subset of S rather than a single element of S. What is essential is to have $R \times S$ rather than S as the basic set of stimuli to which the axioms of Definition 6 apply.

For example, the pair (r_i, σ_j) must be sampled *and* conditioned as a pattern, and the axioms are formulated to require that what is sampled and conditioned be a subset of the presentation set T on a given trial. In this connection to simplify notation I shall often write $T_n = (r_{i,n-1}, \sigma_{j,n})$ rather than

$$T = \{(r_i, \sigma_j)\},$$

but the meaning is clear. T_n is the presentation set consisting of the single pattern (or element) made up of response r_i on trial $n - 1$ and stimulus element σ_j on trial n, and from Axiom S1 we know that the pattern is sampled because it is the only one presented.

From a psychological standpoint something needs to be said about part of the presentation set being the previous response. In the first place, and perhaps most importantly, this is not an *ad hoc* idea adopted just for the purposes of this paper. It has already been used in a number of experimental studies unconnected with automata theory. Several worked-out examples are to be found in various chapters of Suppes and Atkinson (1960).

Secondly, and more importantly, the use of $R \times S$ is formally convenient, but is not at all necessary. The classical S–R tradition of analysis suggests a formally equivalent, but psychologically more realistic approach. Each response r produces a stimulus σ_r, or more generally, a set of stimuli. Assuming again, for formal simplicity just one stimulus element σ_r, rather than a set of stimuli, we may replace R by the set of stimuli S_R, with the purely contingent presentation schedule

$$P(\sigma_{r,n} \quad r_{n-1}) = 1,$$

and in the model we now consider the Cartesian product $S_R \times S$ rather than $R \times S$. Within this framework the important point about the presentation set on each trial is that one component is purely subject-controlled and the other purely experimenter-controlled—if we use familiar experimental distinctions. The explicit use of S_R rather than R promises to be important in training animals to perform like automata, because the external introduction of σ_r reduces directly and significantly the memory load on the animal.[3] The importance of S_R for models of children's language learning is less clear.

DEFINITION 8. *Let* $\mathscr{S} = \langle R \times S, R, E, \mu, X, P\rangle$ *be a stimulus–response model where*

$$R = \{r_0, ..., r_{p-1}\},$$
$$S = \{\sigma_0, ..., \sigma_m\},$$
$$E = \{e_0, ..., e_{p-1}\},$$

and $\mu(S')$ *is the cardinality of* S' *for* $S' \subseteq S$. *Then* $\mathscr{S}$ *asymptotically becomes the automaton* $\mathfrak{A}(\mathscr{S}) = \langle R, S - \{\sigma_0\}, M, r_0, F\rangle$ *if and only if*

(i) *as* $n \to \infty$ *the probability is* 1 *that the presentation set* T_n *is* $(r_{i,n-1}, \sigma_{j,n})$ *for some* i *and* j,

(ii) $M(r_i, \sigma_j) = r_k$ *if and only if* $\lim_{n\to\infty} P(r_{k,n} \mid T_n = (r_{i,n-1}, \sigma_{j,n})) = 1$ *for* $0 \leqslant i \leqslant p-1$ *and* $1 \leqslant j \leqslant m$,

(iii) $\lim_{n\to\infty} P(r_{0,n} \mid T_n = (r_{i,n-1}, \sigma_{0,n})) = 1$ *for* $0 \leqslant i \leqslant p-1$,

(iv) $F \subseteq R$.

A minor but clarifying point about this definition is that the stimulus σ_0 is not part of the alphabet of the automaton $\mathfrak{A}(\mathscr{S})$, because a stimulus is needed to put the automaton in the initial state r_0, and from the standpoint of the theory being worked out here, this requires a stimulus to which the organism will give response r_0.[4] That stimulus is σ_0. The definition also requires that asymptotically the stimulus–response model $\mathscr{S}$ is nothing but the automaton $\mathfrak{A}(\mathscr{S})$. It should be clear that a much weaker and more general definition is possible. The automaton $\mathfrak{A}(\mathscr{S})$ could merely be embedded asymptotically in $\mathscr{S}$ and be only a part of the activities of $\mathscr{S}$. The simplest way to achieve this generalization is to make the alphabet of the automaton only a proper subset of $S - \{\sigma_0\}$ and correspondingly for the responses that make up the internal states of the automaton; they need be only a proper subset of the full set R of responses. This generalization will not be pursued here, although something of the sort will be necessary to give an adequate stimulus–response account of the semantical aspects of language.

[3] Exactly such a procedure has proved very successful in the pigeon experiments with Diebold and Ebbesen mentioned in Footnote 2.

[4] Two other points about the definition are the following. First, in this definition and throughout the rest of the article e_0 is the reinforcement of response r_0, and not the null reinforcement, as in Sect. 3. This conflict arises from the different notational conventions in mathematical learning theory and automata theory. Second, strictly speaking I should write $(\mathfrak{A}_F \mathscr{S})$ because F is not uniquely determined by $\mathscr{S}$.

REPRESENTATION THEOREM FOR FINITE AUTOMATA. *Given any connected finite automaton, there is a stimulus–response model that asymptotically becomes isomorphic to it. Moreover, the stimulus-response model may have all responses initially unconditioned.*

Proof. Let $\mathfrak{A} = \langle A, \Sigma, M, s_0, F\rangle$ be any connected finite automaton. As indicated already, we represent the set A of internal states by the set R of responses; we shall use the natural correspondence s_i–r_i, for $0 \leqslant i \leqslant p - 1$, where p is the number of states. We represent the alphabet Σ by the set of stimuli $\sigma_1, \ldots, \sigma_m$, and, for reasons already made explicit, we augment this set of stimuli by σ_0, to obtain

$$S = \{\sigma_0, \sigma_1, \ldots, \sigma_m\}.$$

For subsequent reference let f be the function defined on $A \cup \Sigma$ that establishes the natural one–one correspondence between A and R, and between Σ and $S - \{s_0\}$. (To avoid some trivial technical points, I shall assume that A and Σ are disjoint.)

We take as the set of reinforcements

$$E = \{e_0, e_1, \ldots, e_{p-1}\},$$

and the measure $\mu(S')$ is the cardinality of S' for $S' \subseteq S$, so that as in Definition 8, we are considering a stimulus–response model $\mathscr{S} = \langle R \times S, R, E, \mu, X, P\rangle$. In order to show that $\mathscr{S}$ asymptotically becomes an automaton, we impose five additional restrictions on $\mathscr{S}$. They are these.

First, in the case of reinforcement e_0 the schedule is this:

$$P(e_{0,n} \mid \sigma_{0,n}) = 1, \tag{1}$$

i.e., if $\sigma_{0,n}$ is part of the presentation set on trial n, then with probability 1 response r_0 is reinforced—note that the reinforcing event $e_{0,n}$ is independent of the actual occurrence of the event $r_{0,n}$.

Second, the remaining reinforcement schedule is defined by the transition table M of the automaton $\mathfrak{A}$. Explicitly, for $j, k \neq 0$ and for all i and n

$$P(e_{k,n} \mid \sigma_{j,n} r_{i,n-1}) = 1 \qquad \text{if and only if} \qquad M(f^{-1}(r_i), f^{-1}(\sigma_j)) = f^{-1}(r_k). \tag{2}$$

Third, essential to the proof is the additional assumption beyond (1) and (2) that the stimuli $\sigma_0, \ldots, \sigma_m$ each have a positive, noncontingent probability of occurrence on each trial (a model with a weaker assumption could be constructed but it is not significant to weaken this requirement). Explicitly, we then assume that for any cylinder set $Y(C_n)$ such that $P(Y(C_n)) > 0$

$$P(\sigma_{i,n}) = P(\sigma_{i,n} \mid Y(C_n)) \geqslant \tau_i > 0 \tag{3}$$

for $0 \leqslant i \leqslant m$ and for all trials n.

Fourth, we assume that the probability ρ_i of response r_i occurring when no conditioned stimuli is sampled is also strictly positive, i.e., for every response r_i

$$\rho_i > 0, \tag{4}$$

which strengthens Axiom R2.

Fifth, for each integer k, $0 \leqslant k \leqslant mp - 1$, we define the set Q_k as the set of states that have exactly k patterns conditioned, and $Qk_{k,n}$ is the event of being in a state that is a member of Q_k on trial n. We assume that at the beginning of trial 1, no patterns are conditioned, i.e.,

$$P(Q_{0,1}) = 1. \tag{5}$$

It is easy to prove that given the sets R, S, E and the cardinality measure μ, there are many different stimulus–response models satisfying restrictions (1)–(5), but for the proof of the theorem it is not necessary to select some distinguished member of the class of models because the argument that follows shows that all the members of the class asymptotically become isomorphic to $\mathfrak{A}$.

The main thing we want to prove is that as $n \to \infty$

$$P(Q_{mp+1,n}) = 1. \tag{6}$$

We first note that if $j < k$ the probability of a transition from Q_k to Q_j is zero, i.e.,

$$P(Q_{j,n} \ Q_{k,n-1}) = 0, \tag{7}$$

moreover,

$$P(Q_{j,n} \mid Q_{k,n-1}) = 0 \tag{8}$$

even if $j > k$ unless $j = k + 1$. In other words, in a single trial, at most one pattern can become conditioned.

To show that asymptotically (6) holds, it will suffice to show that there is an $\epsilon > 0$ such that on each trial n for $0 \leqslant k \leqslant mp < n$ if $P(Q_{k,n}) > 0$,

$$P(Q_{k+1,n+1} \mid Q_{k,n}) \geqslant \epsilon. \tag{9}$$

To establish (9) we need to show that there is a probability of at least ϵ of a stimulus pattern that is unconditioned at the beginning of trial n becoming conditioning on that trial. The argument given will be a uniform one that holds for any unconditioned pattern. Let $r^*\sigma^*$ be such a pattern on trial n.

Now it is well known that for a connected automaton, for every internal state s, there is a tape x such that

$$M(s_0, x) = s \tag{10}$$

and the length of x is not greater than the number of internal states. In terms of stimulus–response theory, x is a finite sequence of length not greater than p of stimulus elements. Thus we may take $x = \sigma_{i_1}, ..., \sigma_{i_p}$ with $\sigma_{i_p} = \sigma^*$. We know by virtue of (3) that

$$\min_{0 \leqslant i \leqslant m} \tau_i = \tau > 0. \tag{11}$$

The required sequence of responses $r_{i_1}, ..., r_{i_p-1}$ will occur either from prior conditioning or if any response is not conditioned to the appropriate pattern, with guessing probability ρ_i. By virtue of (4)

$$\min_{0 \leqslant i \leqslant p-1} \rho_i = \rho > 0. \tag{12}$$

To show that the pattern $r^*\sigma^*$ has a positive probability ϵ, of being conditioning on trial n, we need only take n large enough for the tape x to be "run," say, $n > p + 1$, and consider the joint probability

$$P^* = P(\sigma_n^*, r_{n-1}^*, \sigma_{i_{p-1}}, n-1, r_{i_{p-2}}, n-2, ..., \sigma_{i_1, n-i_p}, r_{0,n-i_p-1}, \sigma_{0,n-i_p-1}). \tag{13}$$

The basic axioms and the assumptions (1)–(5) determine a lower bound on P^* independent of n. First we note that for each of the stimulus element $\sigma_0, \sigma_{i_1}, ..., \sigma^*$, by virtue of (3) and (11)

$$P(\sigma_n^* \mid \cdots) \geqslant \tau, ..., P(\sigma_{0,n-i_p-1}) \geqslant \tau.$$

Similarly, from (4) and (12), as well as the response axioms, we know that for each of the responses $r_0, r_{i_1}, ..., r^*$

$$P(r_{n-1}^* \mid \cdots) \geqslant \rho, ..., P(r_{0,n-i_p-1} \mid \sigma_{0,n-i_p-1}) \geqslant \rho.$$

Thus we know that

$$P^* \geqslant \rho^p \tau^{p+1},$$

and given the occurrence of the event $\sigma_n^* r_{n-1}^*$, the probability of conditioning is c, whence we may take

$$\epsilon = c\rho^p\tau^{p+1} > 0,$$

which establishes (9) and completes the proof.

Given the theorem just proved there are several significant corollaries whose proofs are almost immediate. The first combines the representation theorem for regular languages with that for finite automata to yield:

COROLLARY ON REGULAR LANGUAGES. *Any regular language is generated by some stimulus–response model at asymptote.*

Once probabilistic considerations are made a fundamental part of the scene, we can in several different ways go beyond the restriction of stimulus–response generated languages to regular languages, but I shall not explore these matters here.

I suspect that many psychologists or philosophers who are willing to accept the sense given here to the reduction of finite automata and regular languages to stimulus–response models will be less happy with the claim that one well-defined sense of the concepts of *intention*, *plan*, and *purpose* can be similarly reduced. However, without any substantial new analysis on my part this can be done by taking advantage of an analysis already made by Miller and Chomsky (1963). The story goes like this. In 1960 Miller, Galanter and Pribram published a provocative book entitled *Plans and the Structure of Behavior*. In this book they severely criticized stimulus–response theories for being able to account for so little of the significant behavior of men and the higher animals. They especially objected to the conditioned reflex as a suitable concept for building up an adequate scientific psychology. It is my impression that a number of cognitively oriented psychologists have felt that the critique of S–R theory in this book is devastating.

As I indicated in the introductory section, I would agree that conditioned reflex *experiments* are indeed far too simple to form an adequate scientific basis for analyzing more complex behavior. This is as hopeless as would be the attempt to derive the theory of differential equations, let us say, from the elementary algebra of sets. Yet the more general theory of sets does encompass in a strict mathematical sense the theory of differential equations.

The same relation may be shown to hold between stimulus–response theory and the theory of plans, insofar as the latter theory has been systematically formulated by Miller and Chomsky.[5] The theory of plans is formulated in terms of tote units ("tote" is an acronym for the cycle test–operate–test–exit). A plan is then defined as a tote hierarchy, which is just a form of oriented graph, and every finite oriented graph may be represented as a finite automaton. So we have the result:

COROLLARY. *Any tote hierarchy in the sense of Miller and Chomsky is isomorphic to some stimulus–response model at asymptote.*

5. REPRESENTATION OF PROBABILISTIC AUTOMATA

From the standpoint of the kind of learning models and experiments characteristic of the general area of what has come to be termed *probability learning*, there is near at hand a straightforward approach to probabilistic automata. It is worth illustrating this approach, but it is perhaps even more desirable to discuss it with some explicitness in order to show why it is not fully satisfactory, indeed for most purposes considerably

[5] After this was written Gordon Bower brought to my attention the article by Millenson (1967) that develops this point informally.

less satisfactory than a less direct approach that follows from the representation of deterministic finite automata already discussed.

The direct approach is dominated by two features: a probabilistic reinforcement schedule and the conditioning of input stimuli rather than response–stimulus patterns. The main simplification that results from these features is that the number of states of conditioning and consequently the number of states in the associated Markov chain is reduced. A two-letter, two-state probabilistic automaton, for example, requires 36 states in the associated Markov chain, rather than 64, as in the deterministic case. We have, as before, the three stimuli σ_0, σ_1, and σ_2 and their conditioning possibilities, 2 for σ_0 as before, but now, in the probabilistic case, 3 for σ_1 and σ_2, and we also need as part of the state, not for purposes of conditioning, but in order to make the response-contingent reinforcement definite, the previous response, which is always either r_1 or r_2. Thus we have $2 \cdot 3 \cdot 3 \cdot 2 = 36$. Construction of the trees to compute transition probabilities for the Markov chain follows closely the logic outlined in the previous section. We may define the probabilistic reinforcement schedule by two equations, the first of which is deterministic and plays exactly the same role as previously:

$$P(e_{1,n} \mid \sigma_{0,n}) = 1,$$

and

$$P(e_{1,n} \mid \sigma_{i,n} r_{j,n-1}) = \pi_{ij},$$

for $1 \leqslant i, j \leqslant 2$.

The fundamental weakness of this setup is that the asymptotic transition table representing the probabilistic automaton only holds in the mean. Even at asymptote the transition values fluctuate from trial to trial depending upon the actual previous reinforcement, not the probabilities π_{ij}. Moreover, the transition table is no longer the transition table of a Markov process. Knowledge of earlier responses and reinforcements will lead to a different transition table, whenever the number of stimuli representing a letter of the input alphabet is greater than one. These matters are well known in the large theoretical literature of probability learning and will not be developed further here.

For most purposes of application it seems natural to think of probabilistic automata as a generalization of deterministic automata intended to handle the problem of errors. A similar consideration of errors after a concept or skill has been learned is common in learning theory. Here is a simple example. In the standard version of all-or-none learning, the organism is in either the unconditioned state (U) or the conditioned state (C). The transition matrix for these states is

	C	U
C	1	0
U	c	$1-c$

and, consistent with the axioms of Sec. 3,

$$P(\text{Correct response} \mid C) = 1 \tag{1}$$

$$P(\text{Correct response} \mid U) = \rho \tag{2}$$

and

$$P(U_1) = 1,$$

i.e., the probability of being in state U on trial 1 is 1. Now by changing (2) to

$$P(\text{Correct response} \mid C) = 1 - \epsilon,$$

for $\epsilon > 0$, we get a model that predicts errors after conditioning has occurred.

Without changing the axioms of Sec. 3 we can incorporate such probabilistic-error considerations into the derivation of a representation theorem for probabilistic automata. One straightforward procedure is to postulate that the pattern sampled on each trial actually consists of N elements, and that in addition M background stimuli common to all trials are sampled, or available for sampling. By specializing further the sampling Axioms S1–S4 and by adjusting the parameters M and N, we can obtain any desired probability ϵ of an error.

Because it seems desirable to develop the formal results with intended application to detailed learning data, I shall not state and prove a representation theorem for probabilistic automata here, but restrict myself to considering one example of applying a probabilistic automaton model to asymptotic performance data. The formal machinery for analyzing learning data will be developed in a subsequent paper.

The example I consider is drawn from arithmetic. For more than 3 years we have been collecting extensive data on the arithmetic performance of elementary-school students, in the context of various projects on computer-assisted instruction in elementary mathematics. Prior to consideration of automaton models, the main tools of analysis have been linear regression models. The dependent variables in these models have been the mean probability of a correct response to an item and the mean success latency. The independent variables have been structural features of items, i.e., arithmetic problems, that may be objectively identified independently of any analysis of response data. Detailed results for such models are to be found in Suppes, Hyman, and Jerman (1967) and Suppes, Jerman, and Brian (1968). The main conceptual weakness of the regression models is that they do not provide an explicit temporal analysis of the steps being taken by a student in solving a problem. They can identify the main variables but not connect these variables in a dynamically meaningful way. In contrast, analysis of the temporal process of problem solution is a natural and integral part of an automaton model.

An example that is typical of the skills and concepts encountered in arithmetic is column addition of two integers. For simplicity I shall consider only problems for which the two given numbers and their sum all have the same number of digits. It will be useful to begin by defining a deterministic automaton that will perform the desired addition by outputting one digit at a time reading from right to left, just as the students are required to do at computer-based teletype terminals. For this purpose it is convenient to modify in inessential ways the earlier definition of an automaton. An automaton will now be defined as a structure $\mathfrak{A} = \langle A, \Sigma_I, \Sigma_0, M, Q, s_0 \rangle$ where A, Σ_I and Σ_0 are nonempty finite sets, with A being the set of internal states as before, Σ_I the input alphabet, and Σ_0 the output alphabet. Also as before, M is the transition function mapping $A \times \Sigma_I$ into A, and s_0 is the initial state. The function Q is the output function mapping $A \times \Sigma_I$ into Σ_0.

For column addition of two integers in standard base-10 representation, an appropriate automaton is the following:

$$A = \{0, 1\},$$
$$\Sigma_I = \{(m, n) : 0 \leqslant m, n \leqslant 9\}$$
$$\Sigma_0 = \{0, 1, \ldots, 9\}$$
$$M(k, (m, n)) = \begin{cases} 0 & \text{if } m + n + k \leqslant 9. \\ 1 & \text{if } m + n + k > 9, \end{cases} \qquad \text{for } k = 0, 1.$$
$$Q(k, (m, n)) = (k + m + n) \bmod 10.$$
$$s_0 = 0.$$

Thus the automaton operates by adding first the ones' column, storing as internal state 0 if there is no carry, 1 if there is a carry, outputting the sum of the ones' column modulus 10, and then moving on to the input of the two tens' column digits, etc. The initial internal state s_0 is 0 because at the beginning of the problem there is no "carry."

For the analysis of student data it is necessary to move from a deterministic to a probabilistic automaton. The number of possible parameters that can be introduced is uninterestingly large. Each transition $M(k, (m, n))$ may be replaced by a probabilistic transition $1 - \epsilon_{k,m,n}$ and $\epsilon_{k,m,n}$, and each output $Q(k(m, n))$, by 10 probabilities for a total of 2200 parameters. Using the sort of linear regression model described above we have found that a fairly good account of student performance data can be obtained by considering two structural variables, C_i, the number of carries in problem item i, and D_i, the number of digits or columns. Let p_i be the mean probability of a correct response on item i and let

$$z_i = \log \frac{1 - p_i}{p_i}.$$

The regression model is then characterized by the equation

$$z_i = \alpha_0 + \alpha_1 C_i + \alpha_2 D_i, \tag{1}$$

and the coefficients α_0, α_1, and α_2 are estimated from the data.

A similar three-parameter automaton model is structurally very natural. First, two parameters, ϵ and η, are introduced according to whether there is a "carry" to the next column.

$$P(M(k, (m, n)) = 0 \quad k + m + n \leqslant 9) = 1 - \epsilon$$

and

$$P(M(k, (m, n)) = 1 \mid k + m + n > 9) = 1 - \eta.$$

In other words, if there is no "carry," the probability of a correct transition is $1 - \epsilon$ and if there is a "carry" the probability of such a transition is $1 - \eta$. The third parameter, γ, is simply the probability of an output error. Conversely, the probability of a correct output is:

$$P(Q(k, (m, n)) = (k + m + n) \bmod 10) = 1 - \gamma.$$

Consider now problem i with C_i carries and D_i digits. If we ignore the probability of two errors leading to a correct response—e.g., a transition error followed by an output error—then the probability of a correct answer is just

$$P(\text{Correct Answer to Problem } i) = (1 - \gamma)^{D_i}(1 - \eta)^{C_i}(1 - \epsilon)^{D_i - C_i - 1}. \tag{2}$$

As already indicated it is important to realize that this equation is an approximation of the "true" probability. However, to compute the exact probability it is necessary to make a definite assumption about how the probability γ of an output error is distributed among the 9 possible wrong responses. A simple and intuitively appealing one-parameter model is the one that arranges the 10 digits on a circle in natural order with 9 next to 0, and then makes the probability of an error j steps to the right or left of the correct response δ^j. For example, if 5 is the correct digit, then the probability of responding 4 is δ, of 3 is δ^2, of 2 is δ^3, of 1 is δ^4, of 0 is δ^5, of 6 is δ, of 7 is δ^2, etc. Thus in terms of the original model

$$\gamma = 2(\delta + \delta^2 + \delta^3 + \delta^4) + \delta^5.$$

Consider now the problem

$$\begin{array}{r} 47 \\ +\ 15. \\ \hline \end{array}$$

Then, where d_i = the ith digit response,

$$P(d_1 = 2) = (1 - \gamma),$$
$$P(d_2 = 6) = (1 - \gamma)(1 - \eta) + \eta\delta.$$

Here the additional term is $\eta\delta$, because if the state entered is 0 rather than 1 when the pair (7,5) is input, the only way of obtaining a correct answer is for 6 to be given as the sum of $0 + 4 + 1$, which has a probability δ. Thus the probability of a correct response to this problem is $(1 - \gamma)[(1 - \gamma)(1 - \eta) + \eta\delta]$. Hereafter we shall ignore the $\eta\delta$ (or $\epsilon\delta$) terms.

Returning to Eq. 2 we may get a direct comparison with the linear regression model defined by Eq. 1, if we take the logarithm of both sides to obtain:

$$\log p_i = D_i \log(1 - \gamma) + C_i \log(1 - \eta) + (D_i - C_i - 1) \log(1 - \epsilon), \qquad (3)$$

and estimate $\log 1 - \gamma$, $\log 1 - \eta$, and $\log 1 - \epsilon$ by regression with the additive constant set equal to zero. We also may use some other approach to estimation such as minimum χ^2 or maximum likelihood. An analytic solution of the standard maximum-likelihood equations is very messy indeed, but the maximum of the likelihood function can be found numerically.

The automaton model naturally suggests a more detailed analysis of the data. Unlike the regression model, the automaton provides an immediate analysis of the digit-by-digit responses. Ignoring the $\epsilon\delta$-type terms, we can in fact find the general maximum-likelihood estimates of γ, ϵ, and η when the response data are given in this more explicit form.

Let there be n digit responses in a block of problems. For $1 \leqslant i \leqslant n$ let $\mathbf{x}_i$ be the random variable that assumes the value 1 if the ith response is correct and 0 otherwise. It is then easy to see that

$$P(\mathbf{x}_i = 1) = \begin{cases} (1 - \gamma) & \text{if } i \text{ is a ones'-column digit} \\ (1 - \gamma)(1 - \epsilon) & \text{if it is not a ones' column and there is no carry to the } i\text{th digit} \\ (1 - \gamma)(1 - \eta) & \text{if there is a carry to the } i\text{th digit,} \end{cases}$$

granted that $\epsilon\delta$-type terms are ignored. Similarly for the same three alternatives

$$P(\mathbf{x}_i = 0) = \begin{cases} \gamma \\ 1 - (1 - \gamma)(1 - \epsilon) \\ 1 - (1 - \gamma)(1 - \eta). \end{cases}$$

So for a string of actual digit responses $x_1, \ldots, x_n$ we can write the likelihood function as:

$$L(x_1, \ldots, x_n) = (1 - \gamma)^a \gamma^b (1 - \epsilon)^c (1 - \eta)^d [1 - (1 - \gamma)(1 - \epsilon)]^e [1 - (1 - \gamma)(1 - \eta)]^f, \qquad (4)$$

where a = number of correct responses, b = number of incorrect responses in the ones' column, c = number of correct responses not in the ones' column when the internal state is 0, d = number of correct responses when the internal state is 1, e = number of incorrect responses not in the ones' column when the internal state is 0, and f = number of incorrect responses when the internal state is 1. (In the model statistical independence of responses is assured by the correction procedure.) It is more convenient to estimate $\gamma' = 1 - \gamma$, $\epsilon' = 1 - \epsilon$, and $\eta' = 1 - \eta$. Making this change, taking the log of both sides of (4) and differentiating with respect to each of the variables, we obtain three equations that determine the maximum-likelihood estimates of γ', ϵ', and η':

$$\frac{\partial L}{\partial \gamma'} = \frac{a}{\gamma'} - \frac{b}{1-\gamma'} - \frac{e\epsilon'}{1-\gamma'\epsilon'} - \frac{f\eta'}{1-\gamma'\eta'} = 0,$$

$$\frac{\partial L}{\partial \epsilon'} = \frac{c}{\epsilon'} - \frac{e\gamma'}{1-\gamma'\epsilon'} = 0,$$

$$\frac{\partial L}{\partial \eta'} = \frac{d}{\eta'} - \frac{f\gamma'}{1-\gamma'\eta'} = 0.$$

Solving these equations, we obtain as estimates:

$$\gamma' = \frac{a-c-d}{a+b-c-d},$$

$$\hat{\epsilon}' = \frac{c(a+b-c-d)}{(c+e)(a-c-d)},$$

$$\hat{\eta}' = \frac{d(a-b-c-d)}{(d+f)(a-c-d)}.$$

The most interesting feature of these estimates is that $\hat{\gamma}'$ is just the ratio of correct responses to total responses in the ones' column. The two equations that yield estimates of ϵ' and η' are especially transparent if they are rewritten:

$$(1-\gamma)(1-\epsilon) = \gamma'\epsilon' = \frac{c}{c+e},$$

$$(1-\gamma)(1-\eta) = \gamma'\eta' = \frac{d}{d+f}.$$

Additional analysis of this example will not be pursued here. I do want to note that the internal states 0 and 1 are easily externalized as oral responses and most teachers do indeed require such externalization at the beginning.

To many readers the sort of probabilistic automaton just analyzed will not seem to be the sort of device required to account for language behavior. Certainly the automata that are adequate to analyze arithmetic are simpler in structure than what is needed even for the language output of a 2-year-old child. On the other hand, I have already begun constructing probabilistic automata that will generate the language output of young children and the preliminary results are far from discouraging. Probabilistic automata and their associated probabilistic grammars seem likely to be the right devices for such language analysis.

REFERENCES

Chomsky, N. Review of B. F. Skinner, *Verbal behavior*. *Language*, 1959, **35**, 26–58.

Chomsky, N. Formal properties of grammars. *In* R. D. Luce, R. R. Bush, and E. Galanter (Eds.), *Handbook of mathematical psychology*. Vol. 2, Ch. 12. New York: Wiley, 1963.

Chomsky, N. *Aspects of the theory of syntax*. Cambridge, Massachusetts: M.I.T. Press, 1965.

Estes, W. K., and Suppes, P. Foundations of statistical learning theory. II. The stimulus sampling model. Technical Report. 26, Institute for Mathematical Studies in the Social Sciences, Stanford Univer., 1959.

Katz, J. J., and Postal, P. M. *An integrated theory of linguistic descriptions*. Cambridge, Massachusetts: M.I.T. Press, 1964.

Kleene, S. C. Representation of events in nerve nets and finite automata. *In* C. E. Shannon and J. McCarthy (Eds.), *Automata studies*. Princeton, New Jersey: Princeton Univer. Press, 1956. Pp. 3–41.

Millenson, J. R. An isomorphism between stimulus–response notation and information processing flow diagrams. *The Psychological Record*, 1967, **17**, 305–319.

Miller, G. A., and Chomsky, N. Finitary models of language users. *In* R. D. Luce, R. R. Bush, and E. Galanter (Eds.), *Handbook of mathematical psychology*. Vol. 2, Ch. 13. New York: Wiley, 1963.

Miller, G. A., Galanter, E., and Pribram, K. H. *Plans and the structure of behavior*. New York: Holt, 1960.

Neimark, E., and Estes, W. K. (Eds.), *Stimulus sampling theory*. San Francisco, California: Holden-Day, 1967.

Rabin, M. O. Probabilistic automata. *Information and Control*, 1963, **6**, 230–245; reprinted *in* E. F. Moore (Ed.), *Sequential machines*. Reading, Massachusetts: Addison-Wesley, 1964. Pp. 98–114.

Rabin, M. O., and Scott, D. Finite automata and their decision problems. *IBM Journal of Research and Development*, 1959, **3**, 114–125; reprinted *in* E. F. Moore (Ed.), *Sequential machines*. Reading, Massachusetts: Addison-Wesley, 1964. Pp. 63–91.

Skinner, B. F. *Verbal behavior*. New York: Appleton, 1957.

Suppes, P., and Atkinson, R. C. *Markov learning models for multiperson interactions*. Stanford: Stanford Univer. Press, 1960.

Suppes, P., Hyman, L., and Jerman, M. Linear structural models for response and latency performance in arithmetic on computer-controlled terminals. *In* J. P. Hill (Ed.), *Minnesota symposia on child psychology*. Minneapolis, Minnesota: Univer. of Minnesota Press, 1967. Pp. 160–200.

Suppes, P., Jerman, M., and Brian, D. *Computer-assisted instruction: Stanford's* 1965–66 *arithmetic program*. New York: Academic Press, 1968.

Received: July 7, 1968

PRINTED IN BRUGES, BELGIUM, BY THE ST. CATHERINE PRESS, LTD.