

Deterministic Autonomous Systems

Arie A. Covrigaru and Robert K. Lindsay

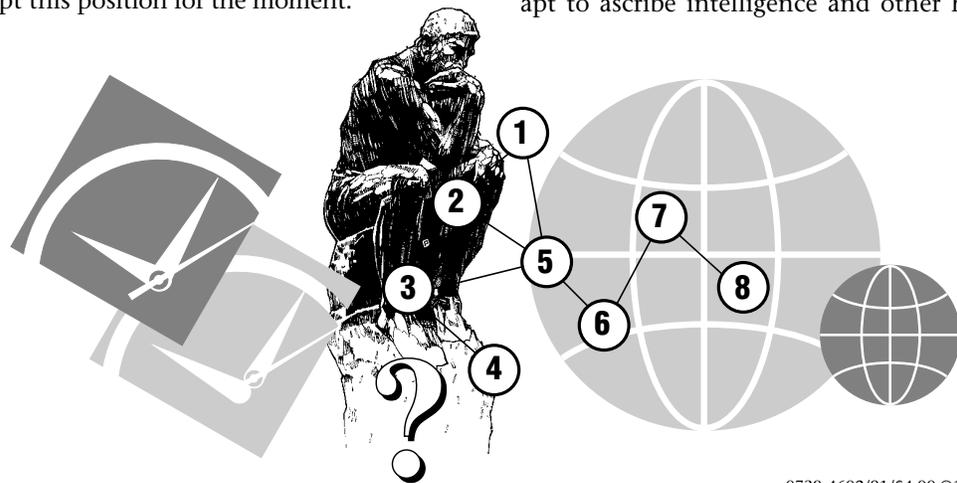
The term intelligence in the phrase artificial intelligence suggests that intelligence is the key characteristic to be analyzed and synthesized by the research discipline. However, for many researchers the objective of this discipline is the scientific understanding of all aspects of complex behavior. For some, this objective might be limited to the traditional goals of scientific psychology: understanding humans. For others, it might include other species and artificial systems. In either case, the enterprise is driven by psychological questions because humans are the extreme of the known range of possibility that drives our curiosity. Therefore, we feel that AI is—or ought to be—seeking to understand and build fully humanlike systems, not simply problem-solving machines. This article presumes that the reader is willing to adopt this position for the moment.

This article argues that autonomy, not problem-solving prowess, is the key property that defines the intuitive notion of “intelligent creature.” To build an intelligent artificial entity that will act autonomously, we must first understand the attributes of a system that lead us to call it autonomous. The presence of these attributes gives autonomous systems the appearance of nondeterminism, but they can all be present in deterministic artifacts and living systems. We argue that autonomy means having the right kinds of goals and the ability to select goals from an existing set, not necessarily creating new goals. We analyze the concept of goals in problem-solving systems in general and establish criteria for the types of goals that characterize autonomy.

Intelligence and Autonomy

Following Turing (1950), most AI researchers accept a purely behavioral criterion for intelligence. That is, an artificial device that is able to do things that are assumed to require intelligence when done by a human merits the description intelligent even though it

is merely a mechanism. Thus, a computer program capable of playing excellent chess would be considered intelligent, even though it succeeds through straightforward, computationally intensive means. The average person is less comfortable with this view. Many would say that a dog or even a worm is intelligent, whereas the chess computer is not, even though a dog, much less a worm, is not capable of anything approaching chess playing. Similarly, the lay person is far more apt to ascribe intelligence and other human



... AI is- or ought to be- seeking to understand and build fully humanlike systems, not simply problem-solving machines.

qualities to a robot than to a computer, even if the robot is capable of only simple actions. What is it that is apparently missing from a chess computer but is apparently present in a robot; a dog; and, perhaps, a worm?

We feel that most would agree that the missing ingredient is autonomy. An entity must be autonomous to be truly intelligent; truly living; and, thus, truly humanoid. Rocks are not autonomous; dogs and, perhaps, worms are. A chess computer, which is intelligent according to Turing's criterion and in the eyes of most computer scientists, is not truly intelligent or alive in the eyes of the lay person because it is not autonomous, in spite of prodigious computational abilities that best all but a handful of people at an intellectual task. However, robots made only approximately in human image and with relatively simple capabilities are readily identified with, empathized with, and anthropomorphized. We suggest this situation is because they seem to have "minds of their own."

Here, the lay person might well claim that the important distinction is between the determinate, predictable behavior of rocks and chess robots and the nondeterminate, unpredictable behavior of living organisms, at least those of sufficient complexity such as the vertebrates. For the scientist, however, this answer is unsatisfactory.

The standard scientific view is that all biological entities, including people, are physical systems that are constrained to obey the laws of physics, and thus, their behavior is mechanistic and deterministic in principle, as is the case for all natural entities and artifacts. There are two qualifications to the deterministic assumption. First, the most successful and fundamental laws of physics, the quantum laws, state that physical processes are inherently nondeterministic; however, quantum indeterminacy is generally assumed to be irrelevant to the macroscopic behavior of animals and computers. Second, the view is also tempered by the recently elucidated phenomena of deterministic, nonrandom but unpredictable behavior known as chaos. These results demonstrate the futility of assuming that the behavior of even simple mechanisms is predictable by closed-form scientific laws,

even if it is the case that the underlying system obeys purely deterministic laws; however, lack of predictability does not conflict with the deterministic assumption but, indeed, makes it more plausible. With these qualifications, we expect that the all-behavior-is-deterministic-in-principle axiom merits the title of standard scientific view, even though doubtless, many scientists, as well as the majority of nonscientists, reject it.

However, if we adopt the standard scientific view, we must account for the apparent autonomy of dogs and the apparent nonautonomy of chess computers not by ascribing indeterminacy to dog behavior or predictability to machine behavior but by some other analysis. Indeterminacy and unpredictability will not suffice as criteria of autonomy.

Why, then, do we ascribe autonomy to some entities and not to others of greater complexity? Clearly, a certain amount of complexity is necessary to merit the description of autonomy, but as the chess computer illustrates, complexity is not the litmus test. Is there one? We argue that an entity is autonomous if it is perceived to have goals, including certain kinds of goals, and is able to select among a variety of goals that it is attempting to achieve. Several other characteristics contribute to the perception of purposiveness. We discuss these in Criteria of Autonomy.

Autonomy and Goals

The major requirement for autonomy is that the entity must be trying to accomplish something; that is, it must be goal directed. Rocks are not considered to be goal directed because it is natural to assume that rocks are merely moving in response to forces and do not care where they end up. There is nothing a rock wants. If our borderline case, the worm, falls short of autonomy in our eyes, it is probably because it is not goal directed. We might be apt to assume that there are things that a worm wants—food, perhaps the proper temperature, perhaps the absence of light, perhaps other worms—but we are not so sure that they are doing more than moving about until they recognize one of these things and then reflexively react to it. We do not need to

The major requirement for autonomy is that the entity must be trying to accomplish something...

assume there is more to a worm, although there might be. Maybe, they are searching, but we cannot tell, so we are not sure if they are autonomous. However, entities that are clearly autonomous force us to adopt the *intentional stance*, that is, to think of them as goal-directed agents.

We know of no way to avoid this subjective criterion; autonomy is in the eye of the beholder. However, we might ask what characteristics of a system description lead us to ascribe autonomy to the system. We attempt to answer this question in the following discussion.

The reader can observe that our straw man, the chess computer, is naturally thought of as goal directed: It is trying to win the game. We propose that what is missing from the chess computer is the right kind of goal.

Indeed, any intelligent system, in particular any problem-solving system, is readily described as having goals that the system is trying to achieve in the course of its operation. Those goals are varied, for example, winning a game of checkers in a checker game program, making the next move in a chess game program, finding the shortest path to a solution, finding the most cost-effective path, learning how to perform better, or learning some new facts. Subgoals are goals that are created, achieved, and abandoned by an intelligent system in the course of achieving higher goals. Some of these goals are preprogrammed in the system by its creators. Some of them are implicit, and others are claimed to be goals independently created by the intelligent component of the system. The question we try to answer here is, What types of goals must a system have to be considered autonomous?

In *Types of Goals in a Problem-Solving System*, we characterize two classes of goals: The first is the class of goals that are non-essential to the definition of autonomy, and the second is the class of goals that are necessary to the concept of autonomy. We believe that the contrast is important to the understanding of autonomy. In each group, we define pairs of types that contrast with each other. When the pair is relevant to autonomy, we show which member of the pair is the one that is required for a system to be considered autonomous. First, however, we explore some other criteria of autonomy.

Criteria of Autonomy

Certain properties of a system tend to elicit in us an intentional stance toward the system. These properties are not individually necessary nor collectively sufficient for autonomy; they

act as evidence. The more of them that are present and the greater their degree (in those cases where they are a matter of degree), the greater our tendency to attribute autonomy to the system.

Prerequisites for Intentionality

Robots, men, dogs, and worms—but not chess machines—are capable of movement, and perhaps, this is one prerequisite of autonomy. Of course, simply being movable is not sufficient. After all, almost everything, including the planets, moves. Movements result from the impingement of energy. However, even rocks fall down when pushed off a cliff, but no one considers them autonomous. The system at least must be able to store energy and use it to initiate movement under appropriate environmental situations, even when the environmental stimulus lacks sufficient energy to directly cause movement. Furthermore, the connection between cause and effect must be adaptive and interactive for the cause to be seen as a stimulus and the effect as a response, that is, as an information-processing event rather than merely a physical event. We try to more clearly spell out the meaning of “adaptive” in this sense.

The strength of evidence is increased if there is some substantial flexibility, variability, and complexity to the repertoire of movements. However, even fairly complex behavioral repertoires, such as possessed by model airplanes, are not sufficient alone. Imagine our chess computer, now augmented with hands, arms, and eyes, that could actually, with internal energy sources such as batteries, move the pieces on the chess board and do so for a wide variety of piece styles and sizes. Such a device would now meet our description of adaptive movement as surely as, say, a fish. We suspect that it would indeed seem much more humanlike than its immobile predecessor, but we also suspect that for most of us, it would fall short of a fish in appearing autonomous.

Fluidity and adaptability of movement would help. Our chess robot, were it capable of deftly grasping a piece between fingers, then rotating its hand to grasp a second piece to place where the removed piece had been, could give our skeptics pause, whereas the same result effected with a jerky pulley-and-gear movement would not pass muster. However, this quality in itself is not an essential quality for autonomy either.

Does the requirement of a varied behavioral repertoire mean that a device or organism that is totally paralyzed by birth or accident could not be considered to be

autonomous and, hence, by the lay person's standard not intelligent? Not necessarily. Suppose there were some way to communicate with it, telepathically say. If these communications had sufficient flexibility of the sort we are discussing, then they could suffice as a surrogate for movement. The key issue seems to be one of flexible, adaptive interaction with the environment. Movement, including perhaps vocalization and gestures, is the usual mode of interaction between living organisms and their environment, but it, too, is not essential. However, if our hypothetical paralyzed organism could in no way interact with the environment, not by movement, eye blink, measurable brain activity, or telepathy, we would probably conclude that it was not autonomous (any more) and was brain dead, even though it maintained metabolism, temperature, and respiration.

An important quality of the interaction is *robustness*, that is, the ability to survive in a variety of situations, a variety so great that not all details could possibly have been anticipated exactly and appropriate responses preplanned. Robustness is required for *self-sufficiency*, the ability to respond in such a way as to avoid danger and remain viable and intact in a varying environment without the intervention of other entities. However, self-sufficiency is a matter of degree itself: Human infants are generally regarded as autonomous (especially by their parents), even though they require frequent intervention to keep them viable in their normal environments.

Perhaps the crucial intuitive core concept of autonomy is independence. That is, an entity is autonomous only if its actions are not controlled or unduly influenced by its environment and other entities. It can go off by itself and function without the assistance or even the awareness of others. What system properties suggest independence?

Independence proves illusive in a deterministic world. If we adopt the standard scientific view, we have a puzzle because deterministic systems, according to the accepted laws of science, are never disentangled from the environment and, indeed, are influenced to some degree by every event, however remote. In fact, independence seems to conflict with the requirement of environmental sensitivity that was discussed earlier. Our answer is that independence must entail selective attention, including the ability to select where and when the input of information and energy are to be effective and when they are not. Again, this condition is not sufficient because most mundane computer programs fit this description. Sometimes pounding your key-

board gets your computer's attention, sometimes it does not.

Another facet of independence is freedom from programming. That is, once the entity exists, further changes to its structure come either entirely from within or by external intervention of a limited kind. The limitation is that the environment can only indirectly impinge on the structure of the entity—by stimulating its fixed set of sense organs. No other entity can get inside and directly make changes. Thus, autonomous entities can learn and can be taught, but they need not be programmed.

To summarize, autonomous systems include a measure of complexity; interaction; movement (preferable fluid); a variety of behaviors; robustness and differential responsiveness to a variety of environmental conditions; selective attention; and independent existence without detailed, knowledgeable intervention. If we add all these criteria together, we are partway home. We claim that a system with all these properties would usually be perceived to be autonomous, provided that in addition it was also seen to be goal directed.

Goals

The term *goal* is usually self-explanatory, but here, there must be a concise basis for what goals are in terms specific to computer programs. Newell (1981) defines a goal to be a body of knowledge of a state of affairs in the environment. He claims that goals are structurally distinguished from other knowledge for them to enter into the behavior of the entity as something that the entity strives to realize. The definition is based on the concept of a problem state space, which is a set of states plus a set of operations that permit attaining one state from another. System activity consists of searching the problem state space by applying the operators.

A goal is a pair

$$G = \{s_c, S_f\},$$

where

S is the set of possible states,

$s_c \in S$ is the current state,

$S_f \subseteq S$ is a set of final states, and

$s_c \notin S_f$ is not one of the final states.

A goal is a description consisting of the system's current state in the state space and a set of states that are called final states. Implicit in the architecture is the intention to get from the current state to one of the states in the set of final states. When the system is in one of these final states, the event marks the termination of this goal. The only constraint on the current state of the system is

that it is not a member of the set of final states as defined for this goal. We consider two goals to be *equivalent* if their sets of final states are identical.

We see later that this definition of goal excludes an important class of goals, namely, those in which the system is in one of the final states and can move among the set of final states, but whose goal is to stay within this set.

Types of Goals in a Problem-Solving System

The goals of a system are often reflected by the name of the system. A program called Chess Player has the goal of playing chess to win (if it doesn't, the sanity of its programmer is in question). At a lower level, the working memory in SOAR (Laird et al. 1990; Laird, Newell, and Rosenbloom 1987), an architecture for a system that is intended to be capable of general intelligence, contains objects of type "goal" that explicitly define goals that the system is working on. Other goals are implicit. For example, the process of learning in an intelligent system has the goal of improving the performance in the future, but this goal might not appear in its list of goals. In the case of problem-solving systems that learn from experience, such as PRODIGY (Minton 1988), SOAR (Laird, Rosenbloom, and Newell 1986), or LEX (Mitchell, Utgoff, and Banerji 1983), even the process of learning is itself a side effect of the performance system. This discussion concentrates on *explicit goals*, those that can be manipulated by a system and are well defined in terms of states in a state space.

The Source of Goals

The first question we ask about our system's goals is where do the goals come from. Are new goals necessary to the characterization of an autonomous system? What is the meaning of saying that a system sets up its own goals? We argue that a system is *autonomous* if it can choose which goal to pursue at any given time from a set of goals, however established. The concept of own goals should not be confused with new goals. Newness is not necessary; if a system has the ability to make a choice, say, between watching TV at home or going to the movies, we could still consider that to be an autonomous decision even if the two goals preexisted in the system.

Every goal-oriented system includes a set of built-in goals that define the functions and behavior of this system. These goals might be embedded in the architecture, for example, as

part of the program in a computer system, or described in the language of an architecture by a set of production rules, as in the OPS system series production language (Forgy 1981). In both cases, these goals exist when the system is created and in most cases do not change with time. These goals must be there if one wants an intelligent system to do anything at all with the knowledge it has or has access to. If a system has no built-in goals, only knowledge, any goal-directed activity will be a reaction to goals set up by some external agent and, thus, are considered goals of the agent and not of the system. A program can also add new goals to its agenda; such goals, established after the system has been created, will be called acquired goals because they come to exist after the system is activated.

Thus, we make the following distinction: *Built-in goals* exist in an intelligent system when the system is "born" and starts its activities. They are part of the definition or program of this particular system and initially direct most of the system's actions. *Acquired goals* are created in a system after the system starts its activities. They are not necessarily subgoals of existing goals within the system in a problem-solving process, yet they are non-existent when the system starts its "life."

There is another dimension of distinction in the source of goals: the way these goals were created. *Endogenous goals* are created by and within the system. These goals might be created as a reaction to some stimulus from the environment or as subgoals in the process of problem solving. *Exogenous goals* are created outside the system and become its goals either at the time the system was designed or through its sensors, but in either case, these are already formulated as goals.

The distinction between endogenous and acquired goals is that endogenous goals are created exclusively through the intelligent mechanisms of the system; that is, the goals are created through learning, problem solving, or interaction with an external environment but not through direct programming. Exogenous goals include the set of built-in goals and also goals that are set by an external agent to the system, for example, by a direct command.

What is the relevance of these types of goals to autonomous systems? It is common to say that autonomous entities create their own goals. However, any multiprocessing operating system such as UNIX, which most would agree is not autonomous, can acquire goals when a new process is started by a user. These goals would fall into the category of

Goals can be distinguished in intelligent systems... and AI systems... according to the period of time the goal is meant to survive.

exogenous, acquired goals: Any AI system, as previously described, has a set of built-in goals to start with. A random generation of goals in a preset domain is also possible. Such goals are unrelated to the other goals of the system, but the system does not have to distinguish between them and others that are considered meaningful goals. Can this mechanism be considered autonomous based only on having endogenous goals? We think not.

Can any system have truly new goals? Certainly, goals can be added to the agenda of an intelligent system by the mechanisms in the system while the system is operating. For new goals to appear in a system, the system needs to either adopt goals created by an agent external to the system in question or create them as subgoals while performing its tasks. Thus, all goals that a system ever has are in one way or another determined by its initial structure. They might be endogenous goals, but they are not new. Therefore, we need to look further to see what is meant when we consider that a system chooses its own goals.

The Hierarchy of Goals

The hierarchy of goals in an intelligent system is important in making this distinction. As noted in the previous subsection, an autonomous entity must have a number of goals to choose from throughout its existence, but the types of goals that are chosen from are also important.

Any nontrivial, goal-directed system has subgoals that derive from the top-level goals. We make this distinction as follows:

Subgoals are created during the process of achieving other goals and depend on the existence of the goals within which they were created. *Top-level goals* are not subgoals. They define the existence of the artificial entity and are independent from other goals in the system. This is not to say that all top-level goals are unrelated, but only that they could be pursued independently and do not stand in a goal-subgoal relation.

The option and capability to make choices among subgoals is part of any problem-solv-

ing process. At each step in the process, the set of subgoals to choose from is determined by the previous step; thus, the subgoals depend on the goal (which is perhaps itself a subgoal) at the previous level in the goal hierarchy. SOAR's architecture, for example, has fully automated the process of generating subgoals. In SOAR's theory, this ability is the essence of the problem-solving process. Still, this ability does not make an entity autonomous. An autonomous entity also needs independent goals that it can choose from. These goals must be independent not only at their own level in the goal hierarchy but also in the set of all existing goals in the system. Thus, the only goals that have relevance to the characterization of an autonomous system are the top-level goals. These goals can be independent of other existing top-level goals and still reside in the same system. Brooks (1987) argues that an autonomous intelligent entity can be decomposed into many decentralized peripheral subsystems, each interacting with the external world. In this context, the interpretation of his argument is that these subsystems are disjoint, top-level goals. Perhaps, the most striking limitation of a chess program is that it has a one-track mind. It cannot make up its mind about what to do because it can only do one thing, namely, play chess. A system can make a choice only if it is presented with more than one option. Thus, an autonomous system would have to have more than one top-level goal.

The Achievement of Goals

Goals can be distinguished in intelligent systems in general and AI systems in particular according to the period of time the goal is meant to survive. To illustrate, consider a common household task (dish washing) that is interpreted in two different ways, each defining a distinct type of goal for this same task.

In the first scenario, the sink is full of dishes. Mother asks her child "Mark, would you please do the dishes?" This goal is relatively easy to achieve. It is set up by the request and achieved

... new goals can only be established by input from the external environment or by the construction of subgoals.

when all the dishes from the sink are clean.

In the second scenario, a butler who just hired a new maid is listing her duties. Among other general instructions, he says, "Make sure there are always clean dishes in the kitchen cabinets." This goal is another type that arises in response to a change in the status quo of the statement "there are clean dishes in the kitchen cabinets." Unlike the previous goal, this one has an unbounded time span, and it is achieved again and again whenever some environment variable goes out of bounds. In essence, the environment is keeping the goal alive, and the task of washing the dishes is one that is terminated over and over.

These two types of goals are called achievable goals and homeostatic goals. *Achievable goals* have a well-defined set of start and final states in the state space; arriving at any one of the final states marks the achievement and termination of such a goal. These goals are the most common type in AI systems. *Homeostatic goals* are achieved continuously. They do not terminate when the system is in one of the final states; when changes occur, meaning the state has changed and is not a final state, activity to reach the final state is reinitiated.

For some homeostatic goals, the process of getting back to the final states defines an achievable goal. We can say that each time the system moves outside the set of final states, it creates an instance of its appropriate achievable goal. However, some homeostatic goals might require that the system always remain in one of the final states (the only viable ones), and thus, some monitoring action must continuously be taken rather than wait for movement out of the set of final states. Such goals do not strictly fit Newell's definition, as given earlier. In some cases, one could modify the definition of goals to include a subset of states called *borderline states*: Only when in a borderline state is action initiated to return to a final state. This solution is only a patch, however; the essential problem is that biological and other physical systems are continuous in time, action, and energy, and Newell's definition attempts to squeeze them into a discrete space. The solution to this problem is not relevant to this article, however; we merely need to note that homeostatic goals, even limited to those that we earlier loosely characterized, are commonly found in systems that are perceived to be autonomous.

Generally, a system that is perceived to be autonomous will have both achievable and homeostatic goals. To exist over time, an autonomous system must have a set of *home-*

ostatic goals. These goals define the internal environment of the system in which it can make choices about achieving its achievable goals. We see the set of homeostatic goals serving as an administrative mechanism or a metalevel to the set of achievable goals that are the real job the system is doing.

The Lifetime of the System

Living systems that are generally regarded as autonomous, such as mammals, seem to be continually engaged in goal-directed activities, at least when they are awake. Achieving one goal is followed by turning to another goal. These goals include but are not limited to homeostatic goals. Of course, to give the appearance of continual goal-driven activities, a system must exist over a period of some time that is long relative to the time it takes to achieve a single nonhomeostatic goal.

Types of Goals Necessary in an Autonomous System

Table 1 summarizes the foregoing distinctions and their relevance to the issue of autonomy. The left column lists pairs of contrasting classes of goals, and the right column points out the class that characterizes autonomy for those classes that were found to be relevant to autonomy. As shown in the table, autonomous systems are characterized by multiple top-level goals, some of which are homeostatic goals. They can have other types of goals, but this capability is irrelevant to their being seen as autonomous.

To summarize, in any deterministic system, autonomous or not, new goals can only be established by input from the external environment or by the construction of subgoals. Autonomy in setting goals is the option to choose the next goal to achieve from a given set of top-level goals. Continually pursuing goals, including homeostatic goals, contributes to the perception of autonomy.

Conclusion

This article tried to clarify the concept of autonomous systems. Autonomy is a subjective property; that is, it is a property of a description of a system, not an objective property of the system that exists independently of observers. Autonomy is a relative, as well as a subjective, property; that is, a system can be more or less autonomous. The charac-

Pairs of contrasting sets of goals	Relevant	Autonomous
Built-in - Acquired	No	
Endogenous - Exogenous	No	
Single - Multiple	Yes	Multiple
Subgoals - Top Level Goals	Yes	Top-level
Achievable - Homeostatic	Yes	Homeostatic

Table 1. This Table Shows the Dimensions Along Which the Types of Goals Are Defined and Their Relevance to the Definition of Autonomy.

teristics listed here are also a matter of degree. Their presence in greater degree enhances the perception of autonomy:

A goal-directed system will be perceived to be autonomous to the degree that (1) it selects tasks (top-level goals) it is to address at any given time; (2) it exists over a period of time that is long relative to the time required to achieve a goal; (3) it is robust, being able to remain viable in a varying environment; (4) some of its goals are homeostatic; (5) there are always goals that are active (instantiated but not achieved); (6) it interacts with its environment in an information-processing mode; (7) it exhibits a variety of complex responses, including fluid, adaptive movements; (8) its attention to stimuli is selective; (9) none of its functions, actions, or decisions need to be fully controlled by an external agent; and (10) once the system starts functioning, it does not need any further programming.

As we characterized it, autonomy does not necessarily entail an extreme level of complexity or intelligence. Simple living organisms are autonomous even though their intelligence (as measured by the complexity of their computational repertoire) is limited. A chess machine, however, can exhibit problem solving but is not autonomous by our criteria. We would also say that it is not intelligent according to conventional wisdom because intelligence, as generally conceived, requires autonomy.

Acknowledgments

We thank John Laird and Yuval Roth for their helpful comments on previous drafts of this article and for useful discussions and advice.

References

Brooks, R. A. 1987. *Intelligence without Representation*, Artificial Intelligence Laboratory,

Massachusetts Institute of Technology.

Forgy, C. L. 1981. OPS5 User's Manual, Computer Science Department, Carnegie Mellon Univ.

Laird, J. E.; Newell, A.; and Rosenbloom, P. S. 1987. SOAR: An Architecture for General Intelligence. *Artificial Intelligence* 33(1): 1-64.

Laird, J. E.; Rosenbloom, P. S.; and Newell, A. 1986. Chunking in SOAR: The Anatomy of a General Learning Mechanism. *Machine Learning* 1(1): 11-46.

Laird, J. E.; Congdon, C. B.; Altmann, E.; and Swedlow, K. 1990. SOAR User's Manual: Version 5.2, Technical Report, CSE-TR-72-90, Univ. of Michigan.

Minton, S. 1988. *Learning Search Control Knowledge: An Explanation-Based Approach*. Boston, Mass.: Kluwer Academic.

Mitchell, T. M.; Utgoff, R. M.; and Banerji, R. 1983. Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics. In *Machine Learning: An Artificial Intelligence Approach*, eds. R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, 163-190. San Mateo, Calif.: Morgan Kaufmann.

Newell, A. 1981. The Knowledge Level, Technical Report, CMU-CS-81-131, Carnegie Mellon Univ.

Turing, A. 1950. Computing Machinery and Intelligence. *Mind* 59:433-460.



Arie Covrigaru is a graduate student in the Ph.D. program in the Computer Science Department at the University of Michigan. He is a member of the SOAR Project Group, with interest in autonomous systems, specifically properties of interrelated tasks in multitasking systems.

Robert Lindsay is a research scientist at the University of Michigan. He received his B.S. and Ph.D. from Carnegie-Mellon University.

