

Is It an Agent, or Just a Program?: A Taxonomy for Autonomous Agents

Stan Franklin and Art Graesser

Institute for Intelligent Systems, University of Memphis, Memphis, TN 38152, USA
stan.franklin@memphis.edu
graesser@cc.memphis.edu

Abstract

The advent of software agents gave rise to much discussion of just what such an agent is, and of how they differ from programs in general. Here we propose a formal definition of an autonomous agent which clearly distinguishes a software agent from just any program. We also offer the beginnings of a natural kinds taxonomy of autonomous agents, and discuss possibilities for further classification. Finally, we discuss subagents and multiagent systems.

1 Introduction

On meeting a friend or colleague that we haven't seen for a while, or a new acquaintance, some version of the following conversation often ensues:

*What are you working on these days?
Control structures for autonomous agents.
Autonomous agents? What do you mean by that?*

A brief explanation is then followed by:

But agents sound just like computer programs. How are they different?

This elicits a more satisfying explanation that distinguishes between agent and program. The nature of this "more satisfying explanation" motivates this essay. After a review of some of the many ways the term "agent" has been used within the context of autonomous agents, we'll propose and defend a notion of autonomous agent that is clearly distinct from a program. This discussion will lead us to a discussion of possible classifications for autonomous agents.

2 What is an agent?

Workers involved in agent research have offered a variety of definitions, each hoping to explicate his or her use of the word "agent." These definitions range from the simple to the lengthy and demanding. We suspect that each of them grew directly out of the set of examples of agents that the definer had in mind. (This is certainly the case for the version we'll propose below.) Let's orient ourselves by examining and comparing some of these definitions.

The MuBot Agent [<http://www.crystaliz.com/logicware/mubot.html>] *"The term agent is used to represent two orthogonal concepts. The first is the agent's ability for autonomous execution. The second is the agent's ability to perform domain oriented reasoning."*

This pointer at definitions come from an online white paper by Sankar Virdhagriswaran of Crystaliz, Inc., defining mobile agent technology. Autonomous execution is clearly central to agency.

The AIMA Agent [17, p. 33] *"An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors."*

AIMA is an acronym for "Artificial Intelligence: a Modern Approach," a remarkably successful new AI textbook that was used in almost 200 colleges and universities in 1995. The authors were interested in software agents embodying AI techniques. Clearly, the AIMA definition depends heavily on what we take as the environment, and on what sensing and acting mean. If we define the environment as whatever provides input and receives output, and take receiving input to be sensing and producing output to be acting, every program is an agent. Thus, if we want to arrive at a useful contrast between agent and program, we must restrict at least some of the notions of environment, sensing and acting.

The Maes Agent [14, p. 108] *"Autonomous agents are computational systems that inhabit some complex dynamic environment, sense and act autonomously in this environment, and by doing so realize a set of goals or tasks for which they are designed."*

Pattie Maes, of MIT's Media Lab, is one of the pioneers of agent research. She adds a crucial element to her definition of an agent: agents must act autonomously so as to "realize a set of goals." Also environments are restricted to being complex and dynamic. It's not clear whether this rules out a payroll program without further restrictions.

The KidSim Agent [18] *"Let us define an agent as a persistent software entity dedicated to a specific purpose. 'Persistent' distinguishes agents from subroutines; agents have their own ideas about how to accomplish tasks, their own agendas. 'Special purpose' distinguishes them from entire multifunction applications; agents are typically much smaller."*

The authors are with Apple. The explicit requirement of persistence is a new and important addition here. Though many agents are "special purpose" we suspect this is not an essential feature of agency.

The Hayes-Roth Agent [7] *"Intelligent agents continuously perform three functions: perception of dynamic conditions in the environment; action to affect conditions in the environment; and reasoning to interpret perceptions, solve problems, draw inferences, and determine actions."*

Barbara Hayes-Roth of Stanford's Knowledge Systems Laboratory insists that agents reason during the process of action selection. If reasoning is interpreted broadly, her agent architecture does allow for reflex actions as well as planned actions.

The IBM Agent [<http://activist.gpl.ibm.com/WhitePaper/ptc2.htm>] *"Intelligent agents are software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user's goals or desires."*

This definition, from IBM's Intelligent Agent Strategy white paper, views an intelligent agent as acting for another, with authority granted by the other. A typical example might be an information gathering agent, though the white paper talks of eight possible applications. Would you stretch "some degree of independence" to include a payroll program? What if it called itself on a certain day of each month?

The Wooldridge–Jennings Agent [20, p. 2] *"... a hardware or (more usually) software-based computer system that enjoys the following properties:*

- *autonomy: agents operate without the direct intervention of humans or others, and have some kind of control over their actions and internal state;*
- *social ability: agents interact with other agents (and possibly humans) via some kind of agent-communication language;*
- *reactivity: agents perceive their environment, (which may be the physical world, a user via a graphical user interface, a collection of other agents, the INTERNET, or perhaps all of these combined), and respond in a timely fashion to changes that occur in it;*
- *pro-activeness: agents do not simply act in response to their environment, they are able to exhibit goal-directed behaviour by taking the initiative."*

The Wooldridge and Jennings definition, in addition to spelling out autonomy, sensing and acting, allows for a broad, but finite, range of environments. They further add a communications requirement. What would be the status of a payroll program with a graphical interface and a decidedly primitive communication language?

The SodaBot Agent [Michael Coen

<http://www.ai.mit.edu/people/sodabot/slideshow/total/P001.html>] *"Software agents are programs that engage in dialogs [and] negotiate and coordinate transfer of information."*

SodaBot is a development environment for software agent being constructed at the MIT AI Lab by Michael Coen. Note the apparently almost empty intersection between this definition and the preceding seven. We say "apparently" since negotiating, for example, requires both sensing and acting. And dialoging requires communication. Still the feeling of this definition is vastly different from the first few, and would seem to rule out almost all standard programs.

The Foner Agent [Lenny Foner — Download from

<ftp://media.mit.edu/pub/Foner/Papers/Julia/Agents--Julia.ps> or online at <http://foner.www.media.mit.edu/people/foner/Julia/> (click on "What's an agent? Crucial notions")]

Foner requires much more of an agent. His agents collaborate with their users to improve the accomplishment of the users' tasks. This requires, in addition to autonomy, that the agent dialog with the user, be trustworthy, and degrade gracefully in the face of a "communications mismatch." However, this quick paraphrase doesn't do justice to Foner's analysis.

The Brustoloni Agent [3, 6 p. 265] *"Autonomous agents are systems capable of autonomous, purposeful action in the real world."*

The Brustoloni agent, unlike the prior agents, must live and act "in the real world." This definition excludes software agents and programs in general. Brustoloni also insists that his agents be *"reactive"* – that is, be able to respond to external, asynchronous stimuli in a timely fashion."

As these definitions make clear, there's no general agreement as to what constitutes an agent, or as to how agents differ from programs. The Software Agents Mailing List on the Internet provides a FAQ (frequently asked questions) that says,

The FAQ Agent [http://www.ee.mcgill.ca:80/~belmarc/agent_faq.html] *"This FAQ will not attempt to provide an authoritative definition ..."*

It does provide a list of attributes often found in agents: Autonomous, goal-oriented, collaborative, flexible, self-starting, temporal continuity, character, communicative, adaptive, mobile, [5]. Several of these would seem to rule out our payroll program.

3 The Essence of Agency

We normally avoid prescriptive arguments about how a word *should* be used. Russell and Norvig put it this way: "The notion of an agent is meant to be a tool for analyzing systems, not an absolute characterization that divides the world into agents and non-agents." [17, p. 33] The only concepts that yield sharp edged categories are mathematical concepts, and they succeed only because they are content free. Agents "live" in the real world (or some world), and real world concepts yield fuzzy categories.

Nevertheless, we will propose a mathematical style definition of an autonomous agent, knowing full well that it must fail around the edges. Our definition attempts to capture the essence of being an autonomous agent, and to define the broadest class of such agents. Further restrictions can then be added to define more particular classes of agents. Ideally, such an endeavor would produce a nomenclature of agents that could be used relatively unambiguously by researchers in the field, resulting in clearer communications.

Note that, though we fall back on the term "agent" as a shorthand, we intend to define the term "autonomous agent." Other people, Luck and D'Inverno [11] for example, first define an agent and then restrict to autonomous agents. We take a different approach.

The definitions of the previous section seem to derive from one or both of two common uses of the word agent: 1) one who acts, or who can act, and 2) one who acts in place of another with permission. Since "one who acts in place of " acts, the second usage requires the first. Hence, let's go for a definition of the first notion.

What are examples of agents in this first sense upon which we can build our mathematical style definition? Well, humans act, as do most *other animals*. (I say most since some animals act during a portion of their lives and not during others, for example the sea squirt [4].) Also, some autonomous mobile robots act, for example Brooks' Herbert [2, p. 8; 6, pp. 263-265]. All of these are real world agents. Software agents "live" in computer operating systems, databases, networks, MUDs, etc. Almost all the definitions in the previous section refer to software agents. Finally, artificial life agents "live" in artificial environments on a computer screen or in its memory [10, 6, pp. 185-208]. What do these agents share that constitutes the essence of being an autonomous agent?

Each is situated in, and is a part on some environment. Each senses its environment and acts autonomously upon it. No other entity is required to feed it input, or to interpret and use its output. Each acts in pursuit of it's own agenda, whether satisfying evolved drives as in humans and animals, or pursuing goals designed in by some other agent, as in software agents. (Artificial life agents may be of either variety.) Each acts so that its current actions may effect its later sensing, that is its actions effect its environment. Finally, each acts continually over some period of time. A software agent, once invoked, typically runs until it decides not to. An artificial life agent often runs until it's eaten or otherwise dies. Of course, some human can pull the plug, but not always. Mobile agents on the Internet may be beyond calling back by the user.

To us, these requirements constitute the essence of being an agent. Let's formalize them into a definition.

*An **autonomous agent** is a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future.*

One way of clarifying the boundaries of this definition is by looking at extreme cases. Humans and some animals are at the high end of being an agent, with multiple, conflicting drives, multiples senses, multiple possible actions, and complex sophisticated control structures (minds [6]) . At the low end, with one or two senses, a single action, and an absurdly simple control structure (mind?) we find a thermostat. A thermostat? Yes, a thermostat satisfies all the requirements of the definition, as does a bacterium. Strange things sometimes happen at the extremes. Espousing a definition entails these risks.

Our definition yields a large and varied class of agents as was to be expected of one requiring only the essence. No doubt it's too large to be useful as is. Adding additional requirements for different purposes will produce useful subclasses of agents. We'll

discuss some of these in a later section. But first, there are a some basic points to clarify.

Autonomous agents are situated in some environment. Change the environment and we may no longer have an agent. A robot with only visual sensors in an environment without light is not an agent. Systems are agents or not with respect to some environment. The AIMA agent discussed above requires that an agent "can be viewed" as sensing and acting in an environment, that is, there must exist an environment in which it is an agent.

What about ordinary programs? A payroll program in a real world environment could be said to sense the world via it's input and act on it via its output, but is not an autonomous agent because its output would not normally effect what it senses later. A payroll program also fails the "over time" test of temporal continuity. It runs once and then goes into a coma, waiting to be called again. Most ordinary programs are ruled out by one or both of these conditions, regardless of how we stretch to define a suitable environment. All software agents are programs, but not all programs are agents.

Nor are software agents defined by their tasks. A spell checker adjunct to a wordprocessor is typically not an autonomous agent for the reasons given in the preceding paragraph. However, a spell checker that watched as I typed and corrected on the fly might well be an autonomous agent. Tasks can be specified so as to require agents to fulfill them.

Subroutines of agents need not be agents for the same reasons that programs need not be. However agents *can* have subagents. Herbert, the robot mentioned above, is built using a subsumption architecture [2], a layered architecture in which each layer senses and acts in order to perform its task. Each layer satisfies all the requirements of an autonomous agent. Thus the layers constitute a multiagent system that controls Herbert. Sumpy [19] is a software agent living in a unix file system. Sumpy, also built using subsumption architecture, consists of subagents that wander, that compress files, that backup files, and that put Sumpy to sleep when the system is busy. Thus, Sumpy is both an agent and a multiagent system.

Our definition above is of an autonomous agent, yet no mention of autonomy appears in the body of the definition. What gives? No explicit mention is needed. An agent that acts in pursuit of its own agenda is acting autonomously. It selects its own actions independently. An extreme case may be illuminating. Consider an agent whose single drive is to act as George tells him to. His modus operandi is to consult George first and then act accordingly. Even this uninteresting agent is acting autonomously in pursuit of his own agenda. His first act, motivated by his drive, is to consult George. His second act, doing what George said to do, is also motivated by this drive. He does what George says, acting autonomously in pursuit of his own agenda. Such an agent seems to be on the fuzzy boundary of our notion of an autonomous agent, and also seems quite useless. Still, it makes a theoretical point.

Wooldridge and Jennings, in their agent definition above, use the term "reactive" to refer to an agents reacting in a timely fashion to changes in its environment. This is

another property that seems to follow from our definition without explicit mention. An agent acting in pursuit of its own agenda will respond to changes in the environment that effect its concerns while ignoring other changes. A successful agent will, presumably, do so in a timely manner.

The term “reactive” is also used in the agent literature in a different sense to refer to actions taken without consulting an internal model. Our definition of autonomous agent allows for purely reactive agents, in this sense, as well as for more deliberative agents that do consult internal models of their environment.

Our definition of an autonomous agents has succeeded in distinguishing between agents and programs. An agent need not be a program at all; it may be a robot or a school teacher. Software agents are, by definition, programs, but a program must measure up to several marks to be an agent. But having given a definition of an autonomous agent, How can we use it to tell if something is one?

4 Describing an Agent

Charles Pettrie, in the panel discussion that followed the presentation of this paper, stressed the need for an operational version of our definition of an autonomous agent. This section is intended as a start in the direction of an operational definition.

If a given system is an autonomous agent, we should be able to describe it in terms of the requirements of the definition given above. This section provides a specification for such a description.

To describe an autonomous agent, describe its

- environment
- sensing capabilities
- actions
- drives
- action selection architecture.

How does one describe an environment? Perhaps using a dynamical systems approach. An environment, E , changes over time and, thus, may be described as a dynamical system $T: X \rightarrow X$, where X is the space of all possible global states of E and T is the global dynamics that updates the current state of E . The state space X is often a vector space, but may take other forms such as a tree with a list at each node [19]. T may be thought of as being an implementation of the physics of E , and may update in either discrete time or in continuous time. Discrete updating may be described by difference equations, and continuous updating by differential equations. T may be deterministic or stochastic (as is our environment based on quantum physics). The updating of E may be influenced by the actions of autonomous agents “living” in the environment. Since each agent living in E is, itself, a part of E , its actions may be thought of as influencing the updating action of T . An initial environment $E(0)$ must be given. $E(t)$ denotes the state of the environment at time t . Each $E(t)$ is an element of the state space X .

If E is a single-agent system, let A denote that agent. *Par abus de language*, let $A:X \rightarrow X$ also denote the updating of E by the actions of the agent A . Let $P:X \rightarrow X$ denote the updating of E other than the actions of A . Together, A and P constitute the global dynamics T of the system.

For a multiagent system, allow finitely many agents A_j , each representing its actions, and P the updating other than these actions.

The sensing capabilities of an autonomous agent A are described in terms of its sensors. Each sensor returns to A some specified portion of the current environment $E(t)$. $E(t)$ is typically a large vector with each attribute of E providing one dimension. Thus the sensing by A , at time t , is given by $S(A,t) = P(E(t))$, where P is a projection onto a subspace depending on the sensing capabilities of A . Sensors are built-in to the agent, though they may degrade over time (as sight and hearing in humans). Sensors may be augmented by the use of instruments (telescopes, jeweler's loops, TV receivers). Sensing may be active, requiring action on the part of the agent, or passive. (All of Sumpster's sensing is active, requiring the issuing of UNIX commands. Human hearing is mostly passive.) Sensing may be internal or external to the agent A . (Proprioception is internal, as is a human being aware of the act of scheduling.)

An action of the agent A produces some change in the current state of its environment. That is, taking the action influences the updating action of T . Describe the action by describing this influence. Each Agent comes with a set of primitive actions built-in. Sequences of these actions can produce higher level actions. In some agents actions can be taken in parallel. The effect of taking a certain action at time t depends on the state of the environment at that time.

Each agent comes with built-in drives (preferences). They evolve in biological agents, and are designed into robotic and computational agents. Some drives may be explicitly represented in the agent's architecture as in Maes' behavior networks [12, 6 pp. 244-255], or in the Ackley-Littman agents [1, 6 pp. 202-206]. Other drives, though not explicitly represented, may be observed from the outside due to causal relations between the functional modules of the agent's architecture. Drives are primitive motivators. Luck and D'Inverno take the existence of such drives as the feature distinguishing autonomous agents from agents [10]. Specific goals are in the service of drives. Agents may have a single drive, or multiple drives. Drives may reinforce or may conflict. The interaction of drives may depend on the state of the agent and on the state of the environment.

An agent's action-selection mechanism decides what to do next and initiates that action. The decision is taken in light of external and internal sensing and in the service of one or more drives. A description of the action-selection mechanism is critical to any description of an agent. Of course, such descriptions can be given at several different levels.

So much for a first cut at an operational definition. Our definition of autonomous agents yields a class of agents so large as not to promise great utility. Let's look at subclasses of agents with more promise.

5 Agent Classifications

The various definitions discussed above involve a host of properties of an agent. Having settled on a much less restrictive definition of an autonomous agent, these properties may help us further classify agents in useful ways. The Table 1 below lists several of the properties mentioned previously.

Agents may be usefully classified according to the subset of these properties that they enjoy. Every autonomous agent, by our definition, satisfies the first four properties. (Note that this does not imply that these four, taken together, are equivalent to our definition.) Adding other properties produces potentially useful classes of agents, for example, mobile, learning agents. Thus a hierarchical classification based on set inclusion occurs naturally. Mobile, learning agents are then a subclass of mobile agents.

Property	Other Names	Meaning
reactive	(sensing and acting)	responds in a timely fashion to changes in the environment
autonomous		exercises control over its own actions
goal-oriented	pro-active purposeful	does not simply act in response to the environment
temporally continuous		is a continuously running process
communicative	socially able	communicates with other agents, perhaps including people
learning	adaptive	changes its behavior based on its previous experience
mobile		able to transport itself from one machine to another
flexible		actions are not scripted
character		believable "personality" and emotional state.

Table 1.

There are, of course, other possible classifying schemes. For example, we might classify software agents according to the tasks they perform, for example, information gathering agents or email filtering agents. Or, we might classify them according to their control architecture. Sumpy, then, would be a fuzzy subsumption agent, while Etzioni and Weld's Softbot would be a planning agent [5]. Agents may also be classified by the range and sensitivity of their senses, or by the range and effectiveness of their actions, or by how much internal state they possess.

Brustoloni's taxonomy of software agents [3] begins with a three-way classification into regulation agents, planning agents, or adaptive agents. A regulation agent,

probably named with regulation of temperature by a thermostat or similar regulation of bodily homeostasis, reacts to each sensory input as it comes in, and always knows what to do. It neither plans nor learns. Planning agents plan, either in the usual AI sense (problem solving agent), or using the case-based paradigm (case-based agents), or using operations research based methods (OR agents), or using various randomizing algorithms (randomizing agent). Brustoloni's adaptive agents not only plan, but learn. Thus there are adaptive problem solving agents, and so on, yielding a two layer taxonomy.

Yet another possible classification scheme might involve the environment in which the agent finds itself, for example software agents as opposed to artificial life agents. And, there must be many, many more such possibilities. Which one, or ones, shall we choose?

6 A Natural Kinds Taxonomy of Agents

In thinking about a taxonomy of agents two possible models come to mind, the biological model and the mathematical model. The biological taxonomy takes the form of a tree with "living creatures" at the root and individual species at the leaves. For example, we humans are classified as

kingdom—animal
 phylum—chordata
 class—mammalia
 order—primate
 family—pongidae
 subfamily—hominidae
 genus—homo
 species—sapiens

where each line represents a branching point of the tree. Might it be possible to create such a taxonomy of autonomous agents? Let's start and see where we get.

At the kingdom level let's classify our agents as either biological, robotic, or computational, as these seem to be natural kinds [9]. Every culture and even very young children readily distinguish between animate organisms, artifacts and abstract concepts. At the phylum level we can reasonably subclassify computational into software agents and artificial life agents. At the class level we might subclassify software agents into task-specific agents (like Sumpy), entertainment agents (like Julia), and computer viruses. At this point we've succeeded in categorizing our major classes of autonomous agents, that is the known families of examples.




Fig. 1. Natural Kinds Classification of Autonomous Agents

7 Further Classification

Suppose we wished to classify software agents further. How might we go about it? The major subclassification schemes that come to mind are via control structures, via environments (database, file system, network, Internet), via language (in which written) or via applications. Each might be useful. Let's try the first.

Let's list some of the possible initial classification schemes for software agents via their control structures. Brustoloni offers regulation, planning and adaptive. Another strategy would be to classify by type of control mechanism, algorithmic, rule-based, planner, fuzzy, neural net, machine learning, etc. Or we might distinguish agents with a central executive from those enjoying distributed control. Other binary classifications might be planning vs. non-planning, learning vs. non-learning, mobile vs. non-mobile, communicative vs. non-communicative, etc.

Suppose we used the binary classification above, including central vs. distributed, in the order mentioned, to create a binary classification tree. The first branching would be according to the first pair. On each of these branches we then branch according to the second pair, and on each of these four we branch again via the third pair, and so on. We've essentially listed a pool of features and classified according to subsets of these features.

Viewing our taxonomic tree from this perspective calls to mind a mathematical taxonomy which also employs collections of properties. A mathematician might define a topological space (please don't bother yourself about the meanings of this mathematical term or others). This essential definition defines the class of spaces to be studied. Then the notion of a Hausdorff space might be defined by an explicit

property of some spaces. Thus the subclass of Hausdorff spaces is specified. Next the notion of a compact space may be defined, yielding the subclass of compact spaces. The intersection of these two is the subclass of compact Hausdorff spaces, about which theorems are often proved. The topological classification continues in this way with defining properties giving rise to subclasses of spaces which are then studied.

This type of classification scheme is known as a matrix organization among psychologists. Each feature defines a dimension. With n features an n -dimensional matrix is created, so that each cell of the matrix corresponds to a collection of features, and provides one possible category for the classification.

Having given the essential definition of an autonomous agent above, the class of autonomous agents is specified. We may then speak of planning agents, or of mobile agents, or even of mobile, communicative, planning agents, each specifying a subclass of agents. Of course, we must have given definitions of these three properties. Having the basic definition of an autonomous agent to build on, and using features for further classification, we may rephrase some of the definitions given earlier in a more convenient manner:

A KidSim Agent is dedicated to a specific purpose, i.e., is a task-specific agent.

A Hayes-Roth Agent reasons to interpret perceptions, solve problems, draw inferences, and determine actions, i.e., is a reasoning agent.

An IBM Agent carries out some set of operations on behalf of a user or another program, i.e., is a task-specific agent.

A Wooldridge-Jennings Agent interacts with other agents (and possibly humans) via some kind of agent-communication language, i.e., is a communicative agent.

A SodaBot Agent engages in dialog, and negotiates and coordinates transfer of information, i.e., is a negotiating, information agent.

8 Subagents and Societies of Agents

Sumpy, the file system maintenance agent mentioned above, can be thought of as a single agent, or as a multiagent system consisting of Wanderer, Compressor, Back-Up and Sleepy. Each of these have independent access to sensors (certain unix commands such as `ls`) and to actions (other unix commands such as `cd`), and each has its own simple agenda. Also, each runs continuously, and acts so as to effect its next sensing. Thus, each may be considered an agent in its own right, and hence a subagent. Sumpy is thus a multiagent system.

Some agents with a layered architecture are not multiagent systems. Müller, Pischel, and Thiel [16] classify such architectures into vertically and horizontally layered. In horizontally layered systems each layer has access to sensing and acting, making a

decomposition into subagents likely. In vertically layered system, only the lowest layer senses, and only the highest acts, making a multiagent decomposition unlikely.

As a multiagent system, Sumpy is particularly simple in that there is almost no communication between the subagents. Each is, of course, privy to sensing initiated by the others, and Sleepy's action effects the others. Also, each subagent sometimes suppresses the actions of the lower layers. One might ask if Wanderer is truly autonomous if Compressor can suppress its actions. A person in jail, or in an elevator, has lost some freedom of movement, but is still autonomous. Environment may be expected to impose limits on an agent's actions.

Going back to our topological analogy, we might call a system with no communication between its subagents a *discrete* multiagent system. A multiagent system in which each agent communicates with every other might be called *fully connected*. Thus multiagents systems can be classified according to the possible communications paths through the system. We might also classify such systems by their communications bandwidth.

In addition to multiagent systems that can reasonably be viewed as constituting a single agent, other multiagent systems are better classified as societies of agents. For example, when a collection of scheduling agents gather to schedule a meeting between their users, they pursue a common goal and intelligent group behavior emerges (see [8] for a similar situation.) Yet, as a group, our definition of agent is not met in that persistence is missing. When scheduling is complete, our agents disperse, perhaps never to gather again in this same grouping. One could argue that the collection of all such scheduling agents at a given site constitute a single agent. To do so, the notions of sensing, acting, and having its own agenda would have to be considerably stretched. As Russell and Norvig have reminded us, the issue here is not truth or falsity, but what's useful in communicating about agents.

The notion of a society of agents leads to a caution. The term "agent" as used by Minsky [15] does not necessarily refer to an autonomous agent as the term is used here. In the context of trying to explain intelligence, Minsky speaks of "mental agents," saying "Each mental agent by itself can only do some simple thing that needs no mind or thought at all." I suspect that some, if not many, of his agents don't meet all our criteria for autonomous agents.

9 Conclusions

An attempt has been made to capture the essence of agency in a formal definition, which allows a clear distinction between a software agent and an arbitrary program. The use of this definition should provide designers of autonomous agents with a checklist of essential features not to be overlooked. In addition, using this definition with further restrictions to useful subclasses should facilitate communications between workers in the field.

We've also offered the beginnings of a theory of autonomous agents as entities in environments described as dynamical systems. This theory, couched as leading to an

operational definition of autonomous agent, should help to clarify the relationship between an agent being designed and its environment, and thus facilitate its design.

The beginnings of a natural kinds taxonomy for autonomous agents is proposed, as is further classification via collections of features. How to continue, or modify this classification could well be the subject of several other papers. There are so many possibilities, each potentially useful in a particular context. Such continuations will surely be pursued.

References

1. David Ackley and Michael Littman. Interactions Between Learning and Evolution. In Christopher Langton et al., ed., *Artificial Life II*. Addison-Wesley 407-509, 1992.
2. Rodney A. Brooks. Elephants Don't Play Chess. In Pattie Maes, ed., *Designing Autonomous Agents*. MIT Press, 1990.
3. Jose C. Brustoloni. Autonomous Agents: Characterization and Requirements. Carnegie Mellon Technical Report CMU-CS-91-204, Carnegie Mellon University 1991.
4. Vincent G. Dethie. The Magic of Metamorphosis: Nature's Own Sleight of Hand. *Smithsonian*, (17): 122 ff, 1986.
5. Oren Etzioni and Daniel Weld. A Softbot-Based Interface to the Internet. *Communications of the ACM*, (37) 7: 72-79, 1994.
6. Stan Franklin. *Artificial Minds*. MIT Press, 1995.
7. B. Hayes-Roth. An Architecture for Adaptive Intelligent Systems. *Artificial Intelligence: Special Issue on Agents and Interactivity*, (72): 329-365, 1995.
8. H. Kautz, B. Selman, and M. Coen. Bottom-up Design of Software Agents. *Communications of the ACM*, (37) 7: 143-146, 1994.
9. F. C. Keil. *Concepts, Kinds, and Cognitive Development*. MIT Press, 1989.
10. Christopher Langton, ed. *Artificial Life*. Addison-Wesley, 1989.
11. Michael Luck and Mark D'Inverno. A Formal Framework for Agency and Autonomy, *Proceedings of the International Conference on Multiagent Systems*. 254-260, 1995.
12. Pattie Maes. How to do the right thing. *Connection Science*. (1):3.1990.
13. Pattie Maes, ed. *Designing Autonomous Agents*, MIT Press, 1991.

14. Pattie Maes. Artificial Life Meets Entertainment: Life like Autonomous Agents. *Communications of the ACM*. (38):11, 108-114, 1995.
15. Marvin Minsky. *The Society of Mind*. Simon and Schuster. 1985.
16. J. P. Müller, M. Pischel, and M. Thiel. Modeling Reactive Behaviour in Vertically Layered Agent Architectures. in Wooldridge and Jennings eds. *Intelligent Agents*. Springer-Verlag, 261-276, 1995.
17. Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
18. D. C. Smith, A. Cypher and J. Spohrer. KidSim: Programming Agents Without a Programming Language. *Communications of the ACM*, (37): 7, 55-67, 1994.
19. Hongjun Song, Stan Franklin and Aregahegn Negatu. SUMPY: A Fuzzy Software Agent. in ed. F. C. Harris, Jr., *Intelligent Systems: Proceedings of the ISCA 5th International Conference* (Reno Nevada, June 1996) International Society for Computers and Their Applications - ISCA, 124-129. 1996.
20. Michael Woolbridge and Nicholas R. Jennings. Agent Theories, Architectures, and Languages: a Survey. in Wooldridge and Jennings eds. *Intelligent Agents*. Springer-Verlag, 1-22, 1995.