

Modeling Agents and their Environment: The Communication Environment

James Odell, James Odell Associates, Ann Arbor, U.S.A.

H. Van Dyke Parunak, Mitchell Fleischer, Altarum, Ann Arbor, U.S.A.

Abstract

Without an environment, an agent is effectively useless. Cut off from the rest of its world, the agent can neither sense nor act. An *environment* provides the conditions under which an entity (agent or object) can exist. It defines the properties of the world in which an agent will function. Designing effective agents requires careful consideration of both the physical and communicational aspects of their environment. This second part of this paper and deals with the communication environment.

1 INTRODUCTION

Agents need to operate and exist within an environment. Figure 1 illustrates a common view that agents perceive their environment through sensors as well as effect actions on it. [Pfeifer, 1999; Weiss, 1999; Russell, 1995] For example, a Stock agent can receive an event indicating that quantities of a particular part are low. The agent then decides whether more parts need to be ordered and, if so, put out a general call-for-proposal so that interested vendors can reply. When proposals arrive, the Stock agent will choose and notify the winning vendor. This model implies that agents interact *via* an environment. Even direct communications (such as vendor notification) must occur through some medium. In other words, the environment provides the appropriate conditions that enable interaction among agents. This insight, largely overlooked in the design of purely electronic agents, is particularly critical for managing agents that are situated in the physical world.




Fig. 1. Agents interact with and through their environment.

2 WHAT IS AN ENVIRONMENT?

One of the key properties of agents is their autonomy. However, autonomy is not an all-or-nothing issue. Practically speaking, agents can neither be totally free of external dependencies nor completely reliant on them. They always depend on external factors to some degree.

An *environment* provides the conditions under which an entity (agent or object) exists.

In other words, it defines the properties of the world in which an agent can and does function. An agent's environment, then, consists not only of all the other entities in its environment, but also those principles and processes under which the agents exist and communicate. Designing effective agents requires careful consideration of all of these factors when designing their environment.

A canonical example of agents situated in an environment is an ant colony. Ants interact with one another largely through chemicals, called *pheromones*, that they deposit in the environment and then sense to guide their actions. Numerous individual interactions yield the emergent development of paths through the environment. However, the environment is more than just a communication channel. Agents depend both on tangible, physical support and on other agents. Two aspects, then, are critical for agent environments (and the formation of paths): the physical and the communicational. The previous paper discussed with the physical environment; this paper discusses with the communication environment.



3 COMMUNICATION ENVIRONMENT

In *individualist* agent environments, agents are viewed as independent entities; whereas in *collectivist* environments, agents are viewed as interdependent. While an agent can operate by alone, the increasing interconnections and networking require a different kind of agent—one that can communicate effectively with other agents. A communication environment provides two things. First, it provides the principles and processes that govern and support the exchange of ideas, knowledge, information, and data. Second, it provides those functions and structures that are commonly employed to enhance communication, such as roles, groups, and the interaction protocols between roles and groups. In short:

The *communication environment* provides those principles, processes, and structures that enable an infrastructure for agents to convey information.

Communication, Interaction, and the Social Agent

Communication

Basically, communication is the conveyance of information from one entity to another. The nature of this transfer can range from the simple to the complex. For example, a satellite could periodically send one bit to inform ground control that it is still functioning correctly; in contrast, the information exchanged within the US Senate to negotiate tax cuts can appear quite chaotic. In contrast, broadcasts such as television commercials do not necessarily result in communication. A signal may go out, but if you are not listening or watching, how can the commercial convey information?

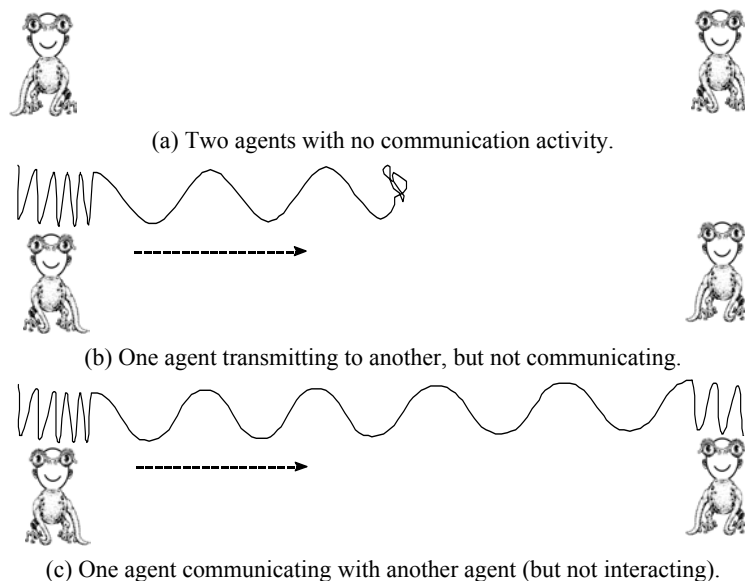


Fig. 2. Agent transmission versus communication.

Figure 2 illustrates the difference between transmission and communication. In Fig. 2(a), neither agent has any transmission activity. Figure 2(b) indicates that the agent on the left transmitted information through the environment, but was not received by the other agent. Communication, however, requires that the information transmitted by one agent results in a state change of another (Fig. 2(c)). In the case of television commercials, perceiving its transmission means that your senses have at least detected it. The perception could involve you buying the advertised goods, throwing a shoe at the television screen, or simply choosing to do nothing. Either way, communication has occurred because the act of sensing and deciding involves a state change by the receiver.

Interaction

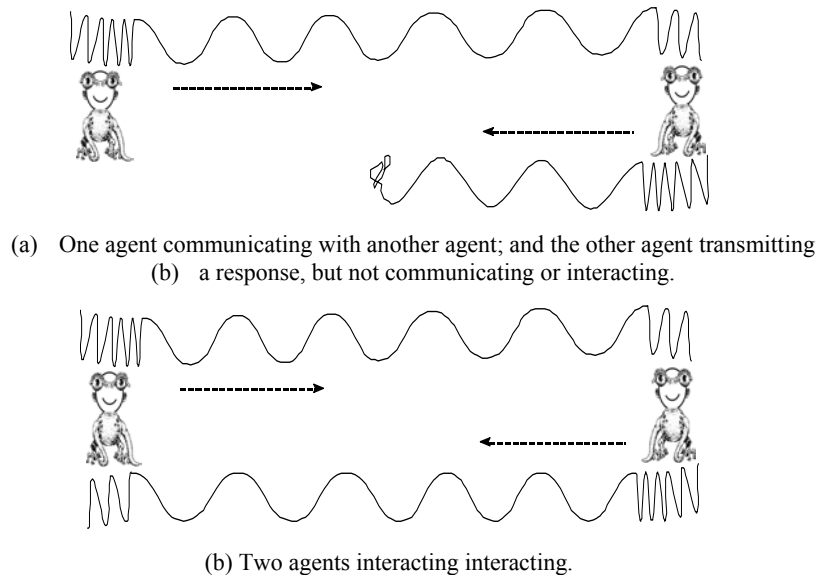



Fig. 3. Agent communication versus interaction.

Proving that communication has occurred, however, requires us to know that the inner state of the receiving agent has in fact changed. We are not advocating that the communication environment possess such mentalistic knowledge—only that such an environment be present so that transmission and communication can occur. However, knowing that a transmission was received can be important to the sending agent. One useful way to determine if communication has occurred is when an interaction results. Figure 3(a) depicts one agent communicating with another. Here, the other agent responds, but the original agent does not receive the responding transmission. (The original agent, then, cannot know for a certainty whether communication occurred.) In other words, there was no interaction between the two agents. *Interaction*¹ requires two-

¹ The action or influence of agents on each other; i.e., having a reciprocal effect.



way communication (i.e., a reciprocal effect), as illustrated in Fig. 3 (b). Interaction, then, not only defines exchange of information, it confirms that the original transmission was in fact received by the other agent. In other words, the original agent can infer that its transmission was communicated to the other agent as soon as a response is received—even if the response communicates only that the responder did not understand the original message.

Social Environment

In agent-based systems, communication and interaction are commonly employed together. Furthermore, agent-based communication can even involve patterns of interaction, or *interaction protocols*. From simply requesting the price of a product to conducting elaborate contract-bidding activities require that some agree-upon approach be in place to facilitate interactive communication—without which the conveyance of information could easily result useless Babel. Such a situation could be considered *social*.

A social environment is a communication environment in which agents interact in a coordinated manner.

As illustrated in Fig. 4, the social environment is a subset the communication environment. In other words, not all communication is social (as defined above), but all social activity requires communication.

Principles of a communication environment

Communication principles provide us with the fundamentals that are essential for interactions, customs, norms, values, commitments, dependencies, and so on, that constitute an agent society. The canonical ant's communication environment is simple: all communications between ants are via pheromones. Here, the communication involves a two-step process: the ant deposits pheromones that act as information for other ants, while the “other ants” query the local environment for the presence of pheromones. In short, these ants participate in a social communication by way of environmental substances, rather than directly with each other. Additionally, some ant societies have multiple kinds of pheromones: one for exploring and one when returning with returning food. The net effect is informing a society of ants about how to find food or home—whichever is useful for any given ant. If an ant is foraging, information about where to find food is useful; if the ant has food, directions on how to get back the colony would be useful. Such interaction is social because it provides the ants with an infrastructure for the colony because it produces coordination among the ants.

Supply-chain agents can have elaborate collaborative protocols acquiring and delivering goods and services along value-adding chains. Defense-related protocols, require different interaction policies at different command levels. Both direct and indirect interaction can be employed as interaction strategies in these applications.





Fig. 4. Social environment: coordination—and possibly cooperation, and competition

In rich multiagent societies (MAS), several principles are required to facilitate the communication environment.

- **Communication language** - Agents communicate to understand and be understood. The formal study of communication has three aspects: syntax, semantics, and pragmatics. Agent-based social environments must define the principles required to address these aspects. Additionally, it must define the types of messages that will be employed (e.g., assertions, queries, replies, requests, and denials) and the ontology. Some of the common agent communication languages (ACL) languages include FIPA ACL, and KQML.
- **Interaction protocols** - An agent interaction protocol (AIP) describes a communication pattern as an allowed sequence of messages between entities and the constraints on the content of those messages [Odell, 2000a and 2000b]. Examples of AIPs include the contract net protocol, Dutch auction protocol, and publish/subscribe protocol. FIPA has standardized more than a dozen AIPs [FIPA, 2000].
- **Coordination strategies** - Agents communicate to achieve their goals and the goals of the social group in which they participate. Cooperation, competition, planning, and negotiation are common principles used to perform activities in a shared environment. AIPs can be associated with each of these strategies.
- **Social policies** - The permissions and obligations that dictate acceptable social behavior. They include being able to apply and enforce these policies across distributed agents and systems. The general focus here is on the application and management of policies on agents and groups of agents—not the detailed management of agent lifecycles and areas currently addressed by FIPA agent management specifications. Other considerations for social policy can involve:
 - Implicit vs explicit rules; not all rules are specified in advance: i.e., learning what the rules are or adjusting to a change in rules, emergence of rules, unconscious rules, when tradition becomes a social norm, or policy.



- Different levels: of influence/power (e.g., the ROI on obeying or violating a rule, strength and “evaporation” of rule; rules don’t always stick around, rules as memes, language use.)
- **Culture** - a set of values, beliefs, desires, intentions, trust, morality. These can determine the characteristics of the above. FIPA vs. KQML cultural differences; English vs. other different-culture language (e.g., Navajo). Culture also affects language, interaction protocol, and social policies (implicit & explicit).

Processes of a communication environment

An agent’s communication environment provides processes that enable agents to interact productively.² In particular, it must provide:

- **Interaction management** - managing the interactions among entities to ensure that they are adhering to the selected agent interaction protocol (AIP). AIP adherence can be maintained by those agents participating in the protocol, so that the environment does not need to be involved. However, trusting that each agent can and will adhere to and ensure correct AIP interaction may not be enough to ensure social order. An environment-level control can be implemented as an AIP-manager agent. Did you get what you wanted/needed/expected;
- **Language processing and policing** - where the language parses correctly, it parses correctly but is wrong (evidence or contradictory), or is correct but inappropriate within the agent’s context.
- **Coordination strategy services**
 - **Directory service** - locating agents can be supported by white-page (individual), yellow-page (industry), or green-page (offered services) methods. In the physical environment, this directory is used to provide information about where the agent is physically; in the social environment, it provides information about an agent’s role or the services that it can provide.
 - **Mediation services** - acting through an intermediate agency. Specialized agents could be established in the environment to act as a communication’s intermediary for activities such as transaction management or ontology translation. Environment-level mediation can be implemented using specialized agents.
- **Policy enforcement service** - control of the agent by its environment or social group. The range of possible mechanisms for enforcing policy mechanisms can range from social sanctions to a complete withdrawal of supporting services for the non-conforming agent.
- **Social differentiation** - the process whereby a group or community becomes separate or distinct. To ensure success, groups will institutionalize and employ roles for their members. An agent can play multiple roles in multiple groups.

² The agent communication *channels* are defined as part of the physical environment. The communication *environment* uses those channels to convey information.

- **Social order** - the production of a structure of relationships among social agents [Castelfranchi, 2000]. Social order can be the result of formal policies as well as emerge via self-organizing mechanisms. The later is a emergent social pattern of its own, such as the stock market. The former has to do with managing the conditions of an agent society as a whole employing a non-accidental and non-chaotic pattern of interactions. For example, auctions employ strict social patterns. Such a mechanism can be employed to control undesirable emergent patterns that need to be remedied. For example, when stock prices rise or fall by too many points in a session, trading curbs are triggered.

Population of a communication environment

As mentioned earlier, a physical environment consists of all those entities in the physical environment. In contrast, an agent's social environment consists of


- those social units in which the agent participates,
- the roles that are employed for social interaction,
- all the other members who play roles in these social units.

Each social unit, or *group*, is a set of agents associated together by some common interest or purpose. There are three reasons for creating groups.

- **Intragroup associations** - Groups are commonly formed to foster or support the interaction of those agents within the group. Here, the group provides a place for a limited number of agents to interact among themselves. For example, such agents might wish to exchange information or seek safety in numbers.
- **Group synergy** - Social units can be formed to take advantage of the synergies of its members, resulting in an entity that enables products and processes that are not possible from any single individual. Corporations, unions, and governments are examples of such social units.
- **Intergroup associations** - Social units also serve as an entity with interactive capability. Here, a group is a set of agents that interact with other sets of agents. Recurrent patterns of interaction define roles, and frequently associated roles are usually considered as defining (sub)groups.

A group can be empty if no agents participate in the group; its collection can also contain a single participating agent or multiple agents. Groups have a separate identity within a larger whole and can be composed of agents, as well as other groups³. Furthermore, groups can become social actors influencing group processes and outcomes, as well. For example, most business organizations interact with sector groups such as industry, technology, agriculture, and government; and each of these can influence the other as well as consist of their own subgroups. In this way, an agent social environment can be thought of as a society where agents interact in a more or less ordered community.

³ Some debate exists about whether a single agent can be its own group, because each agent can be thought of as having both a social and physical existence. There is another debate about whether or not a group has the status of an agent (holonics vs. AALAADIN).



A *role* is an abstract representation of an agent's function, service, or identification within a group. In other words, each role is a class of agents that participates in pattern of dependencies and interactions in a prescribed manner. A pattern of dependencies is an important component of a role. For example, if agent A is a customer, there must be some agent B on whom A depends for goods and services, while B depends on A for money. For AIPs, roles define which actions are permitted for a certain class of agents. For example, an agent playing the customer role may request goods, but not supply them; the supplier has the opposite requirements. [Parunak, 2001]

4 SPATIAL AND TEMPORAL CONSIDERATIONS

An agent's environment—physical or social—must occupy both space and time. Agent populations abide and interact, their processes occur, and their environmental principles are defined over that same temporal space. Agent space and time involves the notion of agent *place*, along with two of its primary attributes: *extent*, and *locality*.

Place

Each agent environment can be thought of as a whole or it can be subdivided into discrete *regions*. Regions partition the agent's physical environment into smaller physical units—where each region may have different or unique characteristics. For example, a grid structure can be defined for the ants so that discrete locations are provided for both the ants and their pheromones, as well as the ability to form pheromone paths. Region definition can also include geographic-based attributes, such as lakes, hills, roads, and structures. In social environments, regions spatially partition the environment into groups and roles. In contrast, temporal space can define unique characteristics for each place in time.

Region specification can include various constraints. For example, in a physical environment we may wish to specify that no two ants may occupy the same place at the same time; yet, we may permit accumulation of multiple units of pheromones. In a social environment, business organizations might be limited to having one person occupying the role of president at any point in time.

Set theoretic distinctions can be made between membership and set. For example, if Agent A belongs to Organization B and Organization B belongs to a federation of organizations C, A does not necessarily belong to C. However, if an Agent A is an element of Set B and Set B is contained by Set C, then Agent A is also an element of Set C.

The region size is determined based on the design granularity: meter-sized places are unrealistic for small ants; micron-sized regions would push the limits of current technology. For example, Pacific Gas & Electric specifies a longitude and latitude within two meters accuracy called a *geocode*. The geocode place size for PG&E, then, is four-square meters. In combat examples, a similar grid structure and size is also employed.

In another example, SRI proposes a new top-level Internet domain called .geo. [SRI, 2000] In a .geo system, the Earth would be partitioned into cells based on latitude and longitude. Dedicated servers would hold the data registered to Web sites within its geographical domain, as well as maps and other information. As illustrated in Fig. 5, places can be arranged hierarchically so that search engines could direct queries to one type of server, depending on what the Web user was looking for. The Internet user could then query for cardiac bench surgery in North America or men's clothing stores in Ann Arbor, Michigan. In this way, web user would never need to use unwieldy .com addresses; the geo-enabled search engine translate a geographic location into web sites registered at that location.

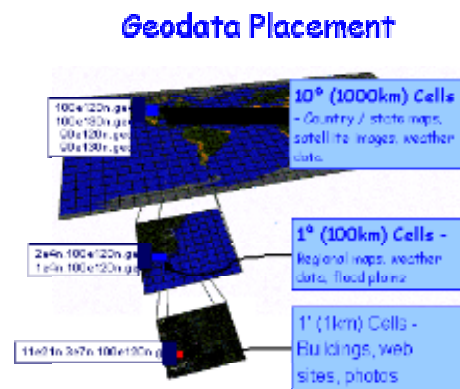



Fig. 5. Places has a hierarchy of geographic placement. [SRI, 2000]

Extent

Agent environments must exist in some designated area (or volume) in space and time. Region designations can be expressed in various ways: length/width/height, location points indicating the boundary, memory or disk locations, to and from dates, and so on. The shape of physical space can also be considered here. Social space can be expressed in terms of the degree of interaction. For example, this could include the number of people you work with, the “degree of separation” between one website and any other website. For example, the environment could be a flat plane or a torus space. In other words, agent environments require an extent that defines its size, shape, and boundaries. Effects of boundary conditions can also be addressed here.

Locality

The ability to locate an entity is an important factor, particularly in agent communication. *Locality* provides the position or situation of a region or entity. Often the locality of the region can become an agent's locality. Locality can be addressed in an absolute and relational manner.



Absolute locators are locators that assign a unique address to each agent or region. Simple two-dimensional grid systems employ column-row designations, geospatial systems employ longitude-latitude-altitude designations, and IT systems employ unique identifiers in the form of keys and unique reference IDs.

In contrast, *relational* locality means that an agent's location can be described as relative to another location. For instance in a connected graph, one agent could be related to other agents, which could in turn be related to other agents. Connected graphs such as the web, electric power networks, or networks of colleagues are examples of where entity's location can be described relative to other entities. In a planar environment an agent's relative neighborhood could be based on physical proximity rather than edges between nodes. For example in a simple two-dimensional grid like a checkerboard, one square can be characterized as diagonal to, or to the side of, and so on. This kind of locality is particularly useful when an entity is constrained to interact with the region of the environment that is near it. For example, ant agents may only move from one region to an adjacent region, and their pheromones might "flow" into neighboring regions where pheromone strength lessens the further it travels.

Locality is useful for several reasons. One primary reason is that communicating with an agent requires that the message can actually be delivered to the agent. The sender of a communication may not be required to know where the receiver is physically located, but at some point the communication service must find the receiver to deliver the message. Another reason is to provide location information. For example, a dispatcher agent might need to know the physical location of its various resources to schedule effectively. Lastly, agent movement or interaction may be based on, or limited to, physical proximity. For example, an ant agent may only move to or interact with the region that is immediately adjacent to it. In contrast, a flea or grasshopper-style agent may jump multiple squares in a single bound, but is limited to a maximum of five.




Fig. 6. Location-aware computing. [Buderi, 2001]

Also, it should be noted that since regions are positioned in an environment, the notion of locality applies to them, as well. Furthermore, an agent's locality can be based on the locality of the region it occupies. For example in the .geo example (Fig. 6), each region had an absolute locator within which other entities are contained. In this way, an

agent's locality can be defined in a discrete space, instead of locating the agent one large, continuous environment.

Hybrid approaches using absolute and relative locality are also useful. AT&T's "bat" transmitters are a good example of a hybrid approach. "Bats" are small battery-powered ultrasonic transmitters that can be worn on a belt or placed inside objects. They broadcast a uniquely identifying 48-bit pulse to receivers embedded every 1.5 meters in ceilings as illustrated in Fig. 6. (For example, about 800 are placed around AT&T's three-story lab in Cambridge, England.) Based on the known position of each receiver, the bearer's precise position can be calculated. In other words, the transmitters and receivers have absolute locators: the transmitter has an 48-bit ID and the receivers are coordinate-based. Then, based on relative proximity, the coordinates of the transmitter can be derived from the receivers' coordinates.

Using this location information, zones can be established around objects and people. If a person's zone overlaps an object's zone, the person becomes the temporary owner of the device, be it a workstation, digital camera, telephone or anything else. There is no logging on and everything the user creates—documents, pictures, memos—is automatically stored in the user's personal files. [Baduri, 2001] In other words this technology, known as *location-aware computing*, detects when you're online and what kind of device you're using. Many companies now have development efforts that involve location-based computing: AT&T's Sentient Computing R&D (described above), IBM's Pervasive Computing Division, HP's CoolTown project, the ubiquitous-computing projects at Intel and Xerox. [Want, 2001] Microsoft is another such company with its new HailStorm services platform. When someone tries to get in touch with you, the HailStorm system will detect your network location and level of accessibility: Are you at your desk? In a meeting? In transit? Depending upon the answer, the system will e-mail, page or call you.

5 POSTLOGUE

By 2015, the social computing is expected by some to morph into ecological or symbiotic computing. John Seely Brown, chief scientist of Xerox suggests that structural matter (atoms) and computing (bits) will become inseparable.

Zillions of sensors, effectors and logical elements (made of organic and inorganic materials) will be interconnected via wireless, peer-to-peer technologies, producing smart, malleable stuff used to build smart appliances, buildings, roads and more. It is during this era that computers disappear. In their place, nearly every physical artifact harbors some computationally based brainpower that helps it know where it was, what was near it, when it was moved and so on. In a way, the inorganic world took on organic properties, using computing to transparently modulate responses to the environment. [Brown, 2001]



6 REFERENCES

- [Brown01] Brown, John Seely, "Where Have All the Computers Gone?," *Technology Review*, 104:1, February 2001, 2001, pp. 86-87.
- [Buderi01] Buderi, Robert (2001) "Computing Goes Everywhere," , 104:1, *Technology Review*, pp. 53-59.
- [Castel00] Castelfranchi, Cristiano (2000) "Engineering Social Order," *Engineering Societies in the Agent World*, Springer, pp. 1-18.
- [FIPA98] FIPA (1998) *Foundation for Intelligent Physical Agents FIPA98 Agent Management Specification*, Geneva, Switzerland, Oct. 1998. (<http://www.fipa.org>)
- [Odell00] Odell, James, H. Van Dyke Parunak, and Bernhard Bauer (2000a) "Representing Agent Interaction Protocols in UML," *First International Workshop on Agent-Oriented Software Engineering 2000*, Paolo Ciancarini and Michael Wooldridge, eds., pp. 58-65.
- [Odell00] Odell, James, H. Van Dyke Parunak, and Bernhard Bauer (2000b) "Extending UML for Agents," *Proc. of the Agent-Oriented Information Systems Workshop at the 17th National Conference on Artificial Intelligence*, Gerd Wagner , Yves Lesperance, and Eric Yu, eds., Austin, TX, pp. 3-17.
- [Parana01] Paranak, H. Van Dyke and James Odell (2001) "Representing Social Structure using UML," *Proc. of the Agent-Oriented Software Engineering Workshop, Agents 2001 conference*, Paolo Ciancarini Michael Wooldridge, and Gerhard Weiss eds., Agents 2001conference, Montreal, Canada, Springer.
- [Pfeife99] Pfeifer, Rolf and Christian Scheier (1999) *Understanding Intelligence*, MIT Press, Cambridge, MA.
- [Russell95] Russell, Stuart and Peter Norvig (1995) *Artificial Intelligence: A Modern Approach*, Prentice-Hall, NJ.
- [SRI00] SRI (2000) http://www.icann.org/tlds/geo1/ModCE/E_TLD_POLICIES_FIN.htm
- [Want01] Want, Roy and Bill Schilit (2001), "Expanding the Horizons of Location-Aware Computing," *IEEE Computer*, August 2001, pp. 31-34.
- [Weiss99] Weiss, Gerhard, ed. (1999) *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, Cambridge, MA.

About the authors



James J. Odell is a consultant, writer, and educator in the areas of object-oriented and agent-based systems, business reengineering, and complex adaptive systems. He has written four books on object orientation and has two books in progress on agent-based system design. His website is www.jamesodell.com.



Dr. Van Parunak is Altarum's Chief Scientist. He is currently working on the applications of complex adaptive systems, with special emphasis on fine-grained agent software architectures for modeling, control, and collaboration. He has written numerous seminal papers in this area, which are available at <http://www.erim.org/~vparunak>.



Dr. Mitchell Fleischer is a Scientific Fellow at Altarum and Practice Area Leader for Supply Chain Engineering. He is an Organizational Psychologist by training and has spent the past 18 years working on issues related to the adoption of new technologies in manufacturing settings. His most recent book is *Concurrent Engineering Effectiveness: Integrating Product Development Across Organizations*, written with Jeffrey Liker. His website is <http://www.erim.org/~mlfleischer>.