

INTELLIGENT AGENTS

Catholijn M. Jonker and Jan Treur

Vrije Universiteit Amsterdam, Department of Artificial Intelligence, Amsterdam, The Netherlands

Keywords: Intelligent agent, Website, Electronic Commerce

Contents

1. Introduction
2. Agent Notions
 - 2.1 Weak Notion of Agent
 - 2.2 Other Notions of Agent
3. Primitive Agent Concepts
 - 3.1 External primitive concepts
 - 3.2 Internal primitive concepts
 - 3.3. An Example Analysis
4. Business Websites
5. A Generic Multi-Agent Architecture for Intelligent Websites
6. Requirements for the Website Agents
 - 6.1 Characteristics and Requirements for the Website Agents
 - 6.2 Characteristics and Requirements for the Personal Assistants
7. The Internal Design of the Information Broker Agents
 - 7.1 A Generic Information Broker Agent Architecture
 - 7.2 The Website Agent: Internal Design
 - 7.3 The Personal Assistant: Internal Design
- Glossary
- Bibliography
- Biographical Sketches

Summary

In this article first an introduction to the basic concepts of intelligent agents is presented. The notion of weak agent is briefly described, and a number of external and internal primitive agent concepts are introduced. Next, as an illustration, application of intelligent agents to business Websites is addressed. An agent-based architecture for intelligent Websites is introduced. Requirements of the agents that play a role in this architecture are discussed. Moreover, their internal design is discussed.

1. Introduction

The term agent has become popular, and has been used for a wide variety of applications, ranging from simple batch jobs and simple email filters, to mobile applications, to intelligent assistants, and to large, open, complex, mission critical systems (such as systems for air traffic control). Some of the key concepts in agent technology lack universally accepted definitions. In particular, there is only partial agreement on what an agent is. For example, simple batch jobs are termed agent because they can be scheduled in advance to perform tasks on a remote machine, mobile

applications are termed agent because they can move themselves from computer to computer, and intelligent assistants are termed agents because they present themselves to human users as believable characters that manifest intentionality and other aspects of a mental state normally attributed only to humans. Besides this variety in different appearances of agents, applications of agents often are concentrated on specific implementations of agents (often in Java). Often the only precise description of an agent is its implementation code, which is dependent on the chosen implementation platform. Therefore, existing agent architectures are often only comparable in an informal manner. A principled design description of an agent at a conceptual and logical level lacks, which makes it difficult to compare agents from different applications.

As agents show a variety of appearances, perform a multitude of tasks, and their abilities vary significantly, attempts have been made to define what they have in common. The *weak notion of agent* is often used as a reference. This notion will be explained in more detail in Section 2; a number of primitive concepts relevant for this type of agent are identified in Section 3. During the design of such agents, these concepts have to be incorporated, and a number of generic agent processes can be identified; for example relating to interaction with the world or to social behaviour with respect to other agents. In the remainder of the paper it is shown how agent technology can be exploited to intelligent Websites to support electronic Commerce in an intelligent manner. The case study of a department store is used as an illustrative example. In Section 4 the application domain of business Websites is discussed. Section 5 introduces an agent-based architecture for an intelligent Website; two types of information agents participating in the application are distinguished. In Section 6 their characteristics and required properties are discussed. In Section 7 a generic information architecture is described and applied to obtain the internal structure of the agents involved in the application.

2. Agent Notions

The *weak notion of agent*, introduced by Wooldridge and Jennings, is often used as a reference in the literature. In this section the notion is explained and illustrated.

2.1 Weak Notion of Agent

The weak notion of agent is a notion that requires the behaviour of agents to exhibit at least the following four types of behaviour:

- Autonomous behaviour
- Responsive behaviour (also called reactive behaviour)
- Pro-active behaviour
- Social behaviour

Autonomy relates to control: although an agent may interact with its environment, the processes performed by an agent are in full control of the agent itself. *Autonomous* behaviour is defined as:

... the system should be able to act without the direct intervention of humans (or other agents) and should have control over its own actions and internal state.

This means that an agent can only be requested to perform some action, and:

The decision about whether to act upon the request lies with the recipient.

Examples of autonomous processes are: process control systems (e.g., thermostats, missile guiding systems, and nuclear reactor control systems), software daemons (e.g., one that monitors a user's incoming email and obtains their attention by displaying an icon when new, incoming email is detected) operating systems.

Many processes that exhibit autonomous behaviour are being termed agents. However, if such agents do not exhibit flexible behaviour, they are not, in general, considered to be intelligent agents. An intelligent agent is defined to be a computer system that is capable of flexible autonomous actions in order to meet its design objectives. Intelligence requires flexibility with respect to autonomous actions, meaning that intelligent agents also exhibit responsive, social, and pro-active behaviour.

An agent exhibits *responsive* (or *reactive*) behaviour if it reacts or responds to new information from its environment. Responsive behaviour is defined as follows:

Agents should perceive their environment (which may be the physical world, a user, a collection of agents, the Internet, etc.) and respond in a timely fashion to changes that occur in it.

A barometer is a simple example of a system that exhibits responsive behaviour: It continually receives new information about the current air pressure and responds to this new information by adjusting its dial.

Pro-active behaviour is defined as follows:

Agents should not simply act in response to their environment, they should be able to exhibit opportunistic, goal-directed behaviour and take the initiative where appropriate.

Pro-active behaviour is the most difficult of the required types of behaviour for an agent defined according to the weak agent notion. For example, pro-active behaviour can occur simultaneously with responsive behaviour. It is possible to respond to incoming new information in an opportunistic manner according to some goals. Also initiatives can be taken in response to incoming new information from the environment, and thus this behaviour resembles responsive behaviour. However, it is also possible to behave pro-actively when no new information is received from the environment. This last behaviour can by no means be called responsive behaviour.

An agent exhibits *social* behaviour if it communicates and co-operates with other agents.

Jennings and Wooldridge define social behaviour as follows:

Agents should be able to interact, when they deem appropriate, with other artificial agents and humans in order to complete their own problem solving and to help others with their activities.

An example of an agent that exhibits social behaviour is a car: it communicates with its human user by way of its dials (outgoing communication dials: speed, amount of fuel, temperature) and its control mechanisms (incoming communication control mechanisms: pedals, the steering wheel, and the gears). It co-operates with its human user, e.g., by going in the direction indicated by the user, with the speed set by that user.

2.2 Other Notions of Agent

Agents can also be required to have additional characteristics. In this section three of these characteristics are discussed: adaptivity, pro-creativity, and intentionality.

Adaptivity is a characteristic that is vital in some systems. An adaptive agent learns and improves with experience. This behaviour is vital in environments that change over time in ways that would make a non-adaptive agent obsolete or give it no chance of survival. This characteristic is modelled often in simulations of societies of small agents, but also, for example, in adaptive user interface agents.

Pro-creativity is of similar importance to find agents that satisfy certain conditions. The chance of survival is often measured in terms of a fitness function. This characteristic is modelled often in simulations of societies of small agents (see the literature in the area of Artificial Life). A computer virus is a very infamous form of a pro-creative agent.

An *intentional system* is defined by Dennett to be an entity ... whose behaviour can be predicted by the method of attributing beliefs, designs and rational acumen.

Mentalistic and intentional notions such as *beliefs, desires, intentions, commitments, goals, plans, preference, choice, awareness*, may be assigned to agents. The *stronger notion of agenthood* in which agents are described in terms of this type of notions provides additional metaphorical support for the design of agents.

3. Primitive Agent Concepts

The notions of agenthood discussed in Section 2 are highly abstract notions. In order to design agents, it is necessary to be familiar with a number of primitive agent concepts. These primitive concepts serve as an ontology or vocabulary used to express analyses and designs of applications of agents and multi-agent systems. Two classes of primitive notions are distinguished: those used to describe the behaviour of agents in terms of their external (or public) states and interactions (Section 3.1), and those used to describe the behaviour of agents in terms of their internal (or private) states, and processes (Section 3.2). In Section 3.3, to illustrate the concepts, an example agent is discussed in terms of these concepts: an elevator.

3.1 External primitive concepts

Two types of interaction of an agent with its environment are distinguished, depending on whether the interaction takes place with an agent or with something else (called an *external world*), for example a database, or the material world. For each of these two types of interaction specific terminology is used.

Interaction with the external world

Two primitive types of interaction with the external world are distinguished. The first type of interaction, *observation*, changes the information the agent has about the world, but does not change the world state itself, whereas the second type, *performing an action*, does change the world state, but does not change the information the agent has about the world. Combinations of these primitive types of interaction are possible; for example, performing an action, and observing its results.

Observation

In which ways is the agent capable of observing or sensing its environment? Two types of observation can be distinguished: the agent passively receives the results of observations without taking any initiative or control to observe (*passive observation*), or the agent actively initiates and controls which observations it wants to perform; this enables the agent to focus its observations and limit the amount of information acquired (*active observation*).

Execution of actions in the external world

An agent may be capable of making changes to the state of its environment by initiating and executing specific types of actions.

Communication with other agents

Two directions of communication are distinguished, which can occur together: *outgoing communication* (is the agent capable of communicating to another agent; to which ones ?), and *incoming communication* (is the agent capable of receiving communication from another agent; from which ones ?).

-
-
-

TO ACCESS ALL THE 20 PAGES OF THIS CHAPTER,
Visit: <http://www.eolss.net/Eolss-sampleAllChapter.aspx>

Bibliography

Brazier, F.M.T., Jonker, C.M., and Treur, J., Principles of Compositional Multi-agent System Development. In: J. Cuenca (ed.), *Proceedings of the 15th IFIP World Computer Congress, WCC'98, Conference on Information Technology and Knowledge Systems, IT&KNOWS'98*, 1998, pp. 347-360. IOS Press. [Paper on DESIRE. More details of the compositional agent system design method DESIRE can be found in them]

Brazier, F.M.T., Jonker, C.M., Treur, J., Compositional Design and Reuse of a Generic Agent Model. *Applied AI Journal*, vol. 14, 2000, pp. 491-538. [Paper on GAM. More details of the generic agent model GAM can be found in them]

Jennings, N.R., and M. Wooldridge (eds.) (1998b), *Agent Technology: Foundations, Applications, and Markets*. Springer Verlag. [Contains papers on basic concepts and application areas of agents.]

Jonker, C.M., Lam, R.A., and Treur, J., A Reusable Multi-Agent Architecture for Active Intelligent Websites. *Journal of Applied Intelligence*, vol. 15, 2001, pp. 7-24. [Paper on information agents and intelligent Websites. More details on the application of information agents to intelligent websites addressed in Sections 4 to 7 can be found in them]

Nwana, H.S., (1996), Software Agents: an Overview, *Knowledge Engineering Review*, vol. 11(3), pp. 205 - 244. [Survey of application areas of agent technology]

Nwana, H.S., and D.T. Ndumu, (1998), A Brief Introduction to Software Agent Technology. In: Jennings, N.R., and M. Wooldridge (eds.) (1998b), *Agent Technology: Foundations, Applications, and Markets*. Springer Verlag., pp. 29 – 47. [Survey of application areas of agent technology]

Wooldridge, M., and Jennings, N.R., Agent theories, architectures, and languages: a survey. In: Wooldridge, M., and Jennings, N.R. (eds.), *Intelligent Agents, Proceedings of the First International Workshop on Agent Theories, Architectures and Languages, ATAL'94*, Lecture Notes in AI, vol. 890, Springer Verlag, 1995, pp. 1-39. [Some of the notions discussed in Section 2 and 3 were adopted from this paper.]

Biographical Sketches

Catholijn Jonker received her Ph.D. degree in Computer Science in 1994 from Utrecht University. Since 1995 she works as assistant professor in the Department of Artificial Intelligence at the Vrije Universiteit Amsterdam. Currently she is coordinator of the Agent Systems Research Group within this department. Her research has focussed on the design and analysis of agent systems and their applications to information agents and Electronic Commerce. Currently the general theme of Catholijn Jonker's research interests is "Dynamics of Behaviour of Agents in a Dynamic Environment." This theme applies to: Multi-issue Negotiation, Intelligent Information Agents, Dynamic Maintenance of Brokering Systems, Requirements Engineering, Verification, Validation, Support systems for RE, Design, and V&V of Agent Systems.

Jan Treur received his Ph.D. in Mathematics and Logic in 1976 from Utrecht University. Since 1986 he works in Artificial Intelligence, from 1990 as a full professor and head of the Department of Artificial Intelligence at the Vrije Universiteit Amsterdam. In the 1990s he headed a research programme on compositional design of knowledge-based systems. In the most recent five years this research programme has focussed on the design and analysis of agent systems and their applications to information agents and Electronic Commerce. Moreover, he addressed the use of agent systems in Biology, Cognitive Science and Social Sciences. From 2001 he has a part-time professorship at the Department of Philosophy in Utrecht. His current research interests include agent systems and their behavioural dynamics, biological, cognitive and social modelling, and philosophy of mind.