

Agent-Oriented Modeling of the Dynamics of Biological Organisms

Catholijn M. Jonker, Jan Treur

Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
URL: <http://www.cs.vu.nl/~{jonker,treur}>, Email: {jonker,treur}@cs.vu.nl

Abstract

In this paper the agent-oriented modeling perspective to cope with biological complexity is discussed. Three levels of dynamics can distinguished and related to each other: dynamics of externally observable agent behavior, dynamics of internal agent processes, and dynamics of multi-agent organisations. This paper addresses the first two. Basic agent concepts to describe externally observable agent behavior are introduced. In the context of two case studies on animal behavior and cell functioning, it is shown how these concepts can be used to specify dynamic properties. In addition, a number of basic agent concepts to describe an agent's internal processes are introduced. Also these concepts are illustrated for specification of dynamic properties in the two case studies. Furthermore, the relationships between dynamic properties of externally observable behavior and dynamic properties of internal agent processes are addressed and illustrated for the animal and cell case studies.

1 Introduction

The area of physical, chemical, or biological processes in nature (or in the laboratory) has the characteristic of involving complex dynamics of multiple active distributed processes and their interaction. This characteristic imposes strong demands on modeling techniques to describe and analyse such processes. Within Computer Science and Artificial Intelligence the last two decennia show a strongly increased activity in modeling complex distributed processes in society and in world wide information and communication networks. Modeling approaches have been developed to cope with this complexity. In particular, the agent-oriented modeling perspective, developed within Artificial Intelligence, has become popular, and has been successfully used for a wide variety of software and hardware applications. Sometimes it is even claimed that agent-orientation is the next trend in Computer Science, as the successor of the object-oriented perspective (Jennings, 2000; Parunak, 1997). Some of the often mentioned reasons of the strong developments and expectations of agent-orientation are that it

- offers a high level modeling perspective, allowing abstraction from unnecessary detail,
- takes multiple active, distributed processes and their interaction into account
- and has an intuitive, almost empathic appeal from the human perspective

In particular, the inherent complexity of the dynamics of multiple, active processes is made manageable by choosing the right level of abstraction in describing them.

The area of physical, chemical, or biological processes in nature has some characteristics in common with the more traditional application domains for agent systems. Therefore, although not

much work has been done in this direction yet, a natural question is whether the agent-oriented modeling perspective is promising for this domain of biological complexity as well. This paper answers this question affirmatively and illustrates the issue by applying and evaluating basic agent modeling concepts and techniques in case studies from Biology.

From the agent-oriented modeling perspective on dynamics, for the single agent case the following levels of modeling can be addressed:

- Modeling the dynamics of *internal processes* within an agent. An example of a well-known agent model is the BDI-model in which Beliefs, Desires and Intentions are used to model the internal dynamics (Rao and Georgeff, 1991; Brazier et al., 1999). Other agent models can be found in (Müller, 1996; Müller, Pischel and Thiel, 1995; Brazier, Jonker and Treur, 2000, 2002)
- Modeling the dynamics of the (*externally observable*) *behavior* of an agent, abstracting from what happens inside the agent. Some approaches can be found in the area of Requirements Engineering for agent behavior; e.g., (Dardenne, Lamsweerde, and Fickas, 1993; Dubois, Du Bois, and Zeippen, 1995; Herlea et al, 1999, 2005)

For each of these levels, within the agent system area modeling languages have been developed that can be used for simulation or for formal analysis of the dynamics. In addition, interlevel relationships between models at different levels have been identified. For example, when a pattern of externally observable behavior is specified (abstracting from internal processes), and a model for internal dynamics is specified, a natural question is whether the model for internal dynamics actually generates the externally observable behavior pattern. This question is directed to the relationship between the descriptions of dynamics at two different aggregation levels (agent-internal and externally observable agent behaviour). For the multi-agent case, addressed in (Bosse, Jonker and Treur, 2006) also for the different levels of aggregation in a multi-agent organisation, comparable interlevel relationships can be identified.

The case studies presented in this paper were chosen in such a way that together they cover the levels pointed out above, and at the same time involve different areas within Biology: animal behavior and cellular processes. The first case study focuses on an *animal* as an agent and models the dynamics of its externally observable (e.g., delayed response or adaptive) behavior and relates these models to high-level models for their internal dynamics. The second case study focuses on the *cell* as an agent, and models the intracellular processes using the BDI-model for the internal functioning of agents.

In this paper, first in Section 2 some characteristics often attributed to agents are discussed. In Section 3 a high level modeling language for the dynamics of agent systems is introduced. After these introductory sections, in Section 4 a number of basic agent concepts to describe externally observable agent behavior are introduced. In the context of the two case studies on animal behavior and cell functioning, it is shown how these concepts can be used to specify dynamic properties. Next, in Section 5 a number of basic agent concepts to describe an agent's internal processes are introduced. Also these concepts are illustrated for specification of dynamic properties in the two case studies mentioned. Section 6 addresses the relationships between dynamic properties of

externally observable behavior and dynamic properties of internal agent processes. Also these relationships are illustrated for the animal and cell case studies. Finally, Section 7 finalises the paper with a discussion.

2 Characteristics of Agents

In an attempt to define in more precise terms when an entity can be considered an agent, the *weak notion of agent* was introduced in (Wooldridge and Jennings, 1995b,c). This is often used as a reference in the literature (see also (Jennings and Wooldridge, 1998a)). The weak notion of agent is a notion that requires the behavior of agents to exhibit the following four types of behavior: autonomous behavior, responsive behavior (sometimes also called reactive behavior), pro-active behavior, and social behavior.

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. This means that an agent can only be requested to perform some action, and, as Jennings and Wooldridge (1998a) state (p. 4): ‘The decision about whether to act upon the request lies with the recipient.’ Examples of autonomous processes are: process control systems (e.g., missile guiding systems, and nuclear reactor control systems), software daemons (e.g., monitoring a user’s incoming email), and operating systems.

Many processes that exhibit autonomous behavior are being termed agents. However, if such agents do not exhibit flexible behavior, they are not, in general, considered to be agents. An agent exhibits *responsive* (or *reactive*) behavior if it perceives its environment and reacts or responds to new information from its environment. A barometer is a simple example of a system that exhibits responsive behavior: It continually receives new information about the current air pressure and responds to this new information by adjusting its dial.

Pro-active behavior is shown as goal-directed behavior or taking the initiative. Pro-active behavior is the most difficult of the required types of behavior for an agent defined according to the weak agent notion. For example, pro-active behavior can occur simultaneously with responsive behavior. It is possible to respond to incoming new information in an opportunistic manner according to some goals. However, it is also possible to behave pro-actively when no new information is received from the environment. A more elaborate treatment of responsive behavior and pro-active behavior and the interactions between them in a co-operation can be found in (Jonker and Treur, 2002).

An agent exhibits *social* behavior if it communicates and co-operates with other agents. An example of an agent that exhibits social behavior 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.

To illustrate these aspects of agent behavior, as an example, living cell behavior is scored for these aspects. The cellular processes take place in an *autonomous* manner, protected from the environment by the cell membrane. The availability of certain substances in the environment affects

choices on which internal processes actually will take place; the cell *responds* to its environment. At the same time it has a built-in initiative to import resources from the environment, and to initiate and maintain metabolic processes and cell division (*pro-activeness*). In cases that other cells are present the cell tunes its behavior to these other cells, using export and import of chemicals to communicate, thus co-operating with them (*social behavior*). This example will be elaborated further in Sections 4.3, 5.5 and 6.3.

3 Specification of Agents and their Dynamic Properties

Dynamics is a major aspect of agents and their behavior: the dynamics of an agent's internal processes, the dynamics of its externally observable behavior, and the joint dynamics of multiple agents in some form of organisation. Therefore, modeling these dynamics is of crucial importance. In this section a high level language is introduced to model and analyse the internal and external dynamics of agents, and of multi-agent organisations, based on the conceptual framework of temporal modelling; cf. (Galton, 2003, 2006).

3.1 State Properties and State Descriptions

To describe dynamics, the notion of state is important. Dynamics will be described in the next section as evolution of states over time. The notion of state as used here is characterised on the basis of an ontology defining a set of physical and/or higher level (state) properties that do or do not hold at a certain point in time. To distinguish these properties from dynamic properties that relate different states over time, these properties are often called *state properties*.

A state property can take the form of variable having a value like 'x has value 5.36', but it also may involve relationships between values of variables, for example 'x < 0.01y', or just qualitative and/or quantitative facts or relationships such as 'the animal A is hungry', 'the animal A has pain', 'the animal A is on top of object B', 'the animal A's body temperature is 37.5° C', 'the environmental temperature is 7° C', 'the color of the animal A equals the color of the environment'. Notice that, depending on the situation, integer or real numbers may or may not be used to describe state properties. Also units may be part of a state property.

A specific state is characterised by a *state description*, which in general is a distinction within the considered set of state properties into the properties that hold in the particular state, and the (other) properties that do not hold in the state. In particular, as a special case, real values for a set of variables are considered possible state descriptions as well. For example, in a dynamical system approach based on variables x_1, x_2, x_3, x_4 , that are related by differential equations over time, values for these variables such as x_1 has value 0.06, x_2 has value 1.84, x_3 has value 3.36, and x_4 has value - 0.27 are considered state descriptions.

State properties are described by ontologies that define the concepts used. To define states and state properties, within agent-oriented analysis the following different types of agent-related ontologies are used. An ontology for *internal properties* of the agent A ($IntOnt(A)$), for properties of the *input* ($InOnt(A)$) and *output* ($OutOnt(A)$) of the agent, and of the (for A) *external* world ($ExtOnt(A)$). All of these ontologies may include assignments of or relationships between real or other numbers, but

also qualitative indications and relationships. For example, properties such as ‘the agent believes that it is raining’, ‘the agent wants an umbrella’, or ‘the agent A has pain’, may belong to $\text{IntOnt}(A)$, whereas ‘the environmental temperature is 7°C ’, may belong to $\text{ExtOnt}(A)$. The agent A’s input ontology $\text{InOnt}(A)$ defines properties for perception, the agent output ontology $\text{OutOnt}(A)$ properties that indicate initiations of actions of A within the external world. The combination of $\text{InOnt}(A)$ and $\text{OutOnt}(A)$ is called the *agent interaction ontology*, defined by $\text{InteractionOnt}(A) = \text{InOnt}(A) \cup \text{OutOnt}(A)$. Note that such an interaction ontology can be based on available ontologies for agent interaction, e.g. (Odell et al., 2000, 2003a, 2003b). The union of all ontologies mentioned above is called the *overall ontology* for A: $\text{OvOnt}(A) = \text{InOnt}(A) \cup \text{IntOnt}(A) \cup \text{OutOnt}(A) \cup \text{ExtOnt}(A)$. If no confusion is expected about the agent to whom ontologies refer, the reference to A is sometimes left out.

To formalise state property descriptions, ontologies are specified in a (many-sorted) first order logical format: an ontology is specified as a finite set of sorts, constants within these sorts, and relations and functions over these sorts (sometimes also called a signature). The example properties mentioned above then can be defined by nullary predicates (or proposition symbols) such as `has_pain`, or by using n-ary predicates (with $n \geq 1$) like `belief(itraining)`, `desire(havinganumbrella)` OR `has_value(x1, 0.06)`, OR `is_of_temperature(environment, 7)`.

For a given ontology Ont , the propositional language signature consisting of all *state ground atoms* based on Ont is denoted by $\text{At}(\text{Ont})$. The *state properties* based on a certain ontology Ont are formalised by the propositions that can be made (using conjunction, negation, disjunction, implication) from the ground atoms and constitute the set $\text{SPROP}(\text{Ont})$.

A *state* for ontology Ont is an assignment of truth-values $\{\text{true}, \text{false}\}$ to the set of ground atoms $\text{At}(\text{Ont})$. The *set of all possible states* for ontology Ont is denoted by $\text{STATES}(\text{Ont})$. In particular, $\text{STATES}(\text{OvOnt})$ denotes the set of all possible *overall states*. For the agent $\text{STATES}(\text{IntOnt})$ is the set of all of its possible *internal states*. Moreover, $\text{STATES}(\text{InteractionOnt})$ denotes the set of all *interaction states*.

3.2 Dynamic Properties

To describe the internal and external dynamics of the agent, explicit reference is made to time. Dynamic properties can be formulated that relate a state at one point in time to a state at another point in time. To express such dynamic properties, and other, more sophisticated ones, the temporal trace language TTL is introduced. This language can be used for any given state ontologies to define dynamic properties based on these state ontologies in a systematic manner.

First, the notion of trace or trajectory is defined more explicitly. Next, the language to express dynamic properties is discussed. A fixed *time frame* τ is assumed which is linearly ordered. Depending on the application, it may be dense (e.g., the real numbers), or discrete (e.g., the set of integers or natural numbers or a finite initial segment of the natural numbers), or any other form, as long as it has a linear ordering.

A *trace* or *trajectory* γ over an ontology Ont and time frame τ is a time-indexed sequence of states $\gamma_t (t \in \tau)$ in $\text{STATES}(\text{Ont})$, i.e., a mapping $\gamma : \tau \rightarrow \text{STATES}(\text{Ont})$. The set of all traces over ontology Ont is denoted by $\text{TRACES}(\text{Ont})$. A *temporal domain description* w is a given set of traces over the

overall ontology, i.e., $W \subseteq \text{TRACES}(\text{OvOnt})$. This set represents all possible developments over time (respecting the world's laws) of the part of the world considered in the application domain.

Given a trace γ over the overall ontology OvOnt , the *input state* of an agent A at time point t , i.e., the restriction $\gamma_t \upharpoonright \text{InOnt}(A)$ of the state of γ at time t to the input ontology is denoted by $\text{state}(\gamma, t, \text{InOnt}(A))$. Analogously, $\text{state}(\gamma, t, \text{OutOnt}(A))$ denotes the *output state* of the agent at time point t , and $\text{state}(\gamma, t, \text{IntOnt}(A))$ the *internal state*. Also here, if no confusion is expected about the particular agent, the reference to A can be left out. We can also refer to the overall state of a system (agent and environment) at a certain moment, denoted by $\text{state}(\gamma, t)$.

To express dynamic properties in a precise manner a language is used in which explicit references can be made to traces, time points and state properties. The sorted predicate logic temporal trace language TTL is built on atoms referring to, e.g., a trace γ , a time point t and a state property p , such as ‘in trace γ at time t for the output state of A property p holds’, formalised by

$$\text{state}(\gamma, t, \text{output}(A)) \models p.$$

This means that the output state property p holds in the state of trace γ at time point t . Here \models is an infix predicate symbol in the language, that can be read as ‘... possesses property ...’. The expression $s \not\models p$ denotes that p does not hold in s (or s does not possess property p). Dynamic properties are expressed by temporal statements built using the usual logical connectives $\&$ (and), \vee (or), \Rightarrow (implies), not (negation) and quantifiers \forall (for all), and \exists (there exists); quantifiers can be used, for example, over traces, time and state properties. A simple example of a dynamic property for any trace γ of an agent A concerns simple reactive behavior (with maximal response time r):

at any point in time t ,
 if in trace γ at t the agent A observes that it is raining,
 then there exists a point in time t' between t and $t + r$ such that at t' the agent A goes to the tree’.

This property expresses that in a certain time interval a response will be shown. As the exact response time is not pre-specified it leaves open a form of nondeterminism in the behavior. This dynamic property is expressed formally in TTL as:

$$\forall t \ [\text{state}(\gamma, t, \text{InOnt}(A)) \models \text{observation_result}(\text{its_raining}) \\ \Rightarrow \exists t' \ t \leq t' \leq t + r \ \& \ \text{state}(\gamma, t', \text{OutOnt}(A)) \models \text{performs}(\text{goto}(\text{tree}))]$$

This shows how dynamic properties of behavior can be mathematically defined without having to give exact values for response time (however, exact upper and lower bounds are given). Alternatively, it could equally well be expressed that exactly at time $t + r$ the animal goes to the tree:

$$\forall t \ [\text{state}(\gamma, t, \text{InOnt}(A)) \models \text{observation_result}(\text{its_raining}) \\ \Rightarrow \text{state}(\gamma, t + r, \text{OutOnt}(A)) \models \text{performs}(\text{goto}(\text{tree}))]$$

Notice that for the sake of presentation this example is kept simple. In the next sections examples of dynamic properties will be presented that are less simple. In particular, as shown in (Jonker and Treur, 2003), it is also possible to model numerical models, for example based on difference and differential equations in TTL.

The Temporal Trace Language TTL can play a useful role in modeling complex phenomena from an agent-oriented perspective in the following manners:

- it provides a way to obtain well-defined and mathematically formalisable *specifications* of *dynamic properties* of externally observable agent behavior, and their internal processes; such dynamic properties can be specified at any level of precision as desired
- for further *analysis* it supports the identification of formalised relationships between different dynamic properties, for example, between properties of an agent's externally observable behavior and its internal processes.
- it offers possibilities to specify and execute *simulation models* in a high level language, for example simulation of an agent's externally observable behavior on the basis of its internal processes.

All of these uses will be illustrated by the case studies and evaluated in Section 7.

3.3 Overview of the Ontology Used to Formalise Agents and their Dynamic Properties

In the sections that will follow an agent ontology will be used to specify relevant aspects and properties of agents. As an overview of the different types of concepts used, Figure 1 is given here. The external agent concepts such as observation, action and communication will be introduced and applied in Section 4. The internal agent concepts will be introduced and applied in Section 5.

4 Modeling Dynamic Phenomena using External Agent Concepts

The description of agents and their behavior discussed in Section 2 is rather vague. In order to model phenomena from an agent-oriented perspective, it is useful to know of a number of basic agent concepts. These agent concepts serve as an ontology or vocabulary used to express conceptualisations and analyses of phenomena based on the agent-oriented perspective. Like all specific modeling perspectives, on the one hand this ontology limits what can be expressed, but on the other hand, it guides the conceptualisation to a form using only concepts that have a body of knowledge and semantics from the domain of agent systems related to them.

Two classes of primitive notions are distinguished: those used to describe the behavior of agents in terms of their externally observable states and behavior in the sense of interactions with their environment (Section 4), and those used to describe the behavior of agents in terms of their internal (or private) states, and processes (addressed in Section 5).

Section 4 addresses basic concepts to model externally observable agent behavior. Such a modeling perspective, abstracts from the internal processes within the agent that realise the behavior, and therefore provides a model at a well-defined more abstract, higher level than the more detailed level of the internal processes. In addition, these concepts provide a vocabulary for the external processes within an agent, which also allows to abstract from the precise physical or chemical realisation of these external states and their dynamics: the concepts are generic over multiple realisations. In Section 4.1 the external agent concepts are introduced, and in Sections 4.2 and 4.3 they are applied to model different (biological) phenomena.

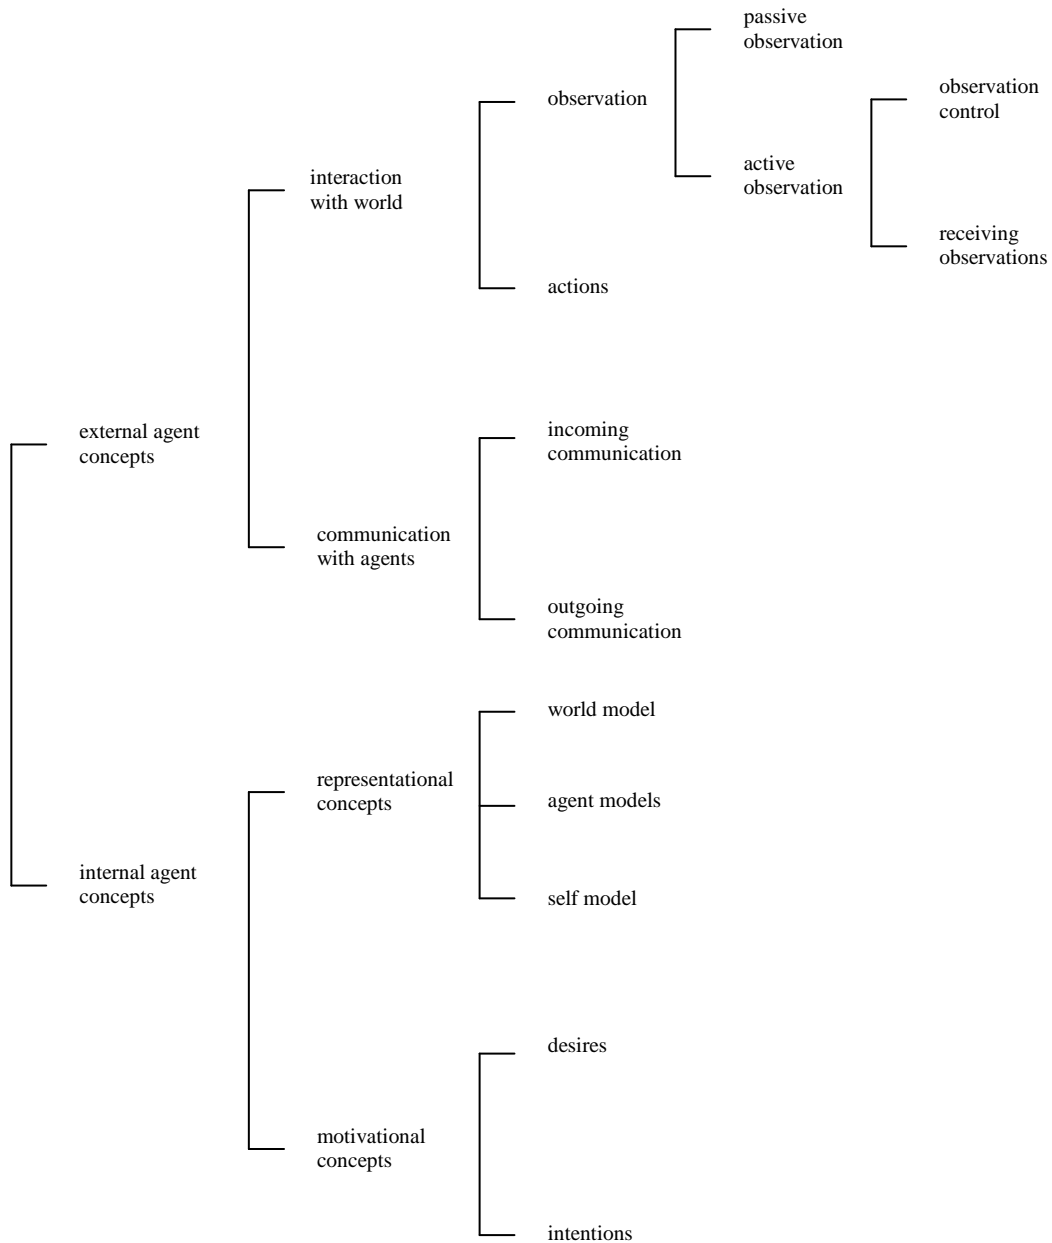


Figure 1 Overview of the types of ontological concepts used

4.1 External Agent Concepts: Interaction with the World and Communication

From the perspective of externally observable behavior, two types of interaction of an agent with its environment are distinguished, depending on whether the interaction takes place with something that is also considered an agent or with something else (called an *external world*). For each of these two types of interaction, specific ontology is used; see also (Odell et al., 2003a, 2003b).

Two primitive types of interaction with the external world are distinguished. The first type of interaction, *observation* (or *sensing*), 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 (directly) 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 performs; this enables the agent to focus its observations and limit the amount of information acquired (*active observation*). An example of the first case is a sensor that generates information all the time, whereas an example of the second case is a sensor that can be controlled in making it active or not, or focusing it.

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

Interaction with parts of the environment that are considered agents as well takes the form of communication. Two directions of communication are distinguished, which can occur together: *outgoing communication* (is the agent capable of communicating to another agent; which information, to which ones?), and *incoming communication* (is the agent capable of receiving communication from another agent; which information, from which ones?).

4.2 Animal Behavior in Terms of External Agent Concepts

One of the examples used in this paper is an animal in an experimental laboratory environment (e.g., a mouse or rat) consisting of a removable screen that separates the animal from positions at which food can be put, which can be covered by a cup. For simplicity the sensing is restricted to visual sensing. In Table 1 an overview is given of how concepts such as observation, action, and communication can be used to describe animal behavior.

Table 1 External Agent Concepts for an Animal

External agent concepts	<i>animal</i>
<i>Interaction with the world</i>	
observations	seeing the presence of food, screen, cup
passive observations	by noticing objects in the visual field without controlling the gaze direction
active observations	by control of gaze direction
performing actions	moving, eating
<i>Communication with other agents</i>	
outgoing communication	uttering sounds
incoming communication	hearing and interpreting sounds uttered by another animal

4.2.1 Dynamic properties for delayed response behavior

For a specific experiment, the animal's externally observed behavior can be described in the following manner (for four different animals). In such a behavior trace the observations and actions of the animal over time are denoted.

Table 2 Example set of observed delayed response traces

<i>time trace</i>	<i>time point 0</i>	<i>time point 1</i>	<i>time point 2</i>	<i>time point 3</i>	<i>time point 4</i>	<i>time point 5</i>
<i>trace 1</i>	food at p2 screen	cup at p2 screen	cup at p2 screen	cup at p2 no screen	cup at p2 no screen goes to p2	cup at p2 no screen
<i>trace 2</i>	no food at p2 screen	no food at p2 screen	cup at p2 screen	cup at p2 screen	cup at p2 no screen	cup at p2 no screen
<i>trace 3</i>	no food screen	no food no screen	no food no screen	no food no screen	food no screen	food no screen goes to p2
<i>trace 4</i>	food at p2 screen	no food at p2 screen	food at p2 screen	cup at p2 screen	cup at p2 no screen	cup at p2 no screen goes to p2

A pattern behind such observed traces (in the literature on animal behavior called 'delayed response behavior') can be expressed in the form of the following dynamic property, where P is a given position:

EDR1

- at any point in time t,
- if in trace γ at t the agent observed that no screen is present
- and at some earlier point in time the agent observed that food was present at position P,
- then for some t' with $t \leq t' \leq t+r$ the agent will go to position P

Here r is the maximal response time. In the formal language of TTL this is written as:

$$\forall t_2 \left[\text{state}(\gamma, t_2, \text{input}(A)) \models \text{observation_result}(\text{no_screen}) \ \& \right. \\ \left. \exists t_1 \leq t_2 \ \text{state}(\gamma, t_1, \text{input}(A)) \models \text{observation_result}(\text{at_position}(\text{food}, P)) \right. \\ \left. \Rightarrow \exists t_3 \ t_2 \leq t_3 \leq t_2+r \ \& \ \text{state}(\gamma, t_3, \text{output}(A)) \models \text{performs}(\text{goto}(P)) \right]$$

Notice that this dynamic property EDR1 (EDR stands for External view on Delayed Response behavior) is formulated as a statement that relates certain state properties over time. These state properties are expressed using the basic external agent concepts, such as observation and action. No references to any internal state properties is made. This makes that such a description of dynamics is generic over multiple realisations of the internal processes. Also the term 'delayed response behavior' (in contrast to, for example, 'shows to have memory') assigned in the literature on animal behavior reflects this independence of internal processes; e.g., (Vauclair, 1996).

This specification forces that in all cases where no cups are present and food was visible but disappeared, after which the absence of food is visible for the animal, still the animal has to go to the position P. An alternative specification for delayed response behavior is the following dynamic property.

EDR2

at any point in time t ,
if in trace γ at some earlier point in time t_1 the agent observed that no screen was present,
and at every point in time t' after t_1 up to t ,
the agent did not observe that a screen was present,
and at some earlier point in time t_2 the agent observed that food was present at P
and at every point in time t'' after t_2 up to t ,
the agent did not observe that no food was present at P ,
then for some t_3 with $t \leq t_3 \leq t+r$ the agent will go to position P

In the formal language of TTL this is written as:

$$\begin{aligned} \forall t \quad & [\exists t_1 \leq t \quad \text{state}(\gamma, t_1, \text{input}(A)) \models \text{observation_result}(\text{not screen_present}) \ \& \\ & \forall t' [t_1 \leq t' \leq t \Rightarrow \text{state}(\gamma, t', \text{input}(A)) \not\models \text{observation_result}(\text{screen_present})] \\ & \& \\ & \exists t_2 \leq t \quad \text{state}(\gamma, t_2, \text{input}(A)) \models \text{observation_result}(\text{at_position}(\text{food}, P)) \ \& \\ & \forall t'' [t_2 \leq t'' \leq t \Rightarrow \text{state}(\gamma, t'', \text{input}(A)) \not\models \text{observation_result}(\text{not at_position}(\text{food}, P))] \\ \Rightarrow & \exists t_3 \quad t \leq t_3 \leq t+r \ \& \ \text{state}(\gamma, t_3, \text{output}(A)) \models \text{performs}(\text{goto}(P))] \end{aligned}$$

This dynamic property EDR2 describes essentially different behavior as compared to EDR1. For example, an animal that behaves according to dynamic property EDR1 will go to position P if it observed food at that position, even if it saw the food disappear immediately (before it was going to P), for example because another animal did eat it. This is because, once the food was observed, the condition of EDR1 on the food remains fulfilled even if the animal observed later that the food is not present anymore. In such a case an animal behaving according to dynamic property EDR2 will not (necessarily) go to P because one of the conditions of EDR2 is not fulfilled, i.e., the condition that in the mean time it was not observed that there is no food at P is not fulfilled. By analysing the behavior in this manner such subtle differences can be identified, expressed and formalised.

4.2.2 Dynamic properties for relative adaptive behavior

This analysis approach can be applied to other types as behavior as well, for example, motivation-based (or goal-directed) behavior or adaptive behavior. Motivation-based behavior will be discussed in Section 4. The following is a simplified example of (relative) adaptive behavior, viewed from an external perspective.

RA (Relative Adaptivity)

for any two traces γ_1 and γ_2 ,
if initially (at time point 0) the skill level in γ_2 is at least as high as in γ_1 ,
and at each time point t
the agent A 's exercising in γ_2 at t is at least as intensive as A 's exercising in γ_1 at t ,
then in trace γ_2 at each point in time t ,
agent A 's demonstrated skill level is at least as high as A 's demonstrated skill level at t in trace γ_1 .

In formalised form (assuming some way of measuring exercising intensity and level of skill):

$\forall \gamma_1, \gamma_2 \in W$

[state(γ_1 , 0, OutOnt(A)) |= has_level(demo_skill, w1) & state(γ_2 , 0, OutOnt(A)) |= has_level(demo_skill, w2) \Rightarrow w1 \leq w2] &
 [$\forall t$ [state(γ_1 , t, InOnt(A)) |= has_intensity(exercising, v1) & state(γ_2 , t, InOnt(A)) |= has_intensity(exercising, v2) \Rightarrow v1 \leq v2]
 $\Rightarrow \forall t$ [state(γ_1 , t, OutOnt(A)) |= has_level(demo_skill, w1) & state(γ_2 , t, OutOnt(A)) |= has_level(demo_skill, w2) \Rightarrow w1 \leq w2]]

4.3 Cell Behavior in Terms of External Agent Concepts

The external agent concepts introduced in Section 4.1 are illustrated by a second example: a cell is analysed from the agent perspective using these basic concepts; see Table 3.

A cell can perform *observations* of the environment by getting specific environmental chemical substances inside via its membrane. This can take place in two manners, passive or active, depending on whether this process is controllable (active) by the cell or not (passive). Passive observation is based on diffusion (which rarely occurs); this process takes place automatically and is not controllable by the cell. Active observation is needed for chemical substances that do not enter the cell by diffusion. To import them, a specific import reaction has to be performed by the cell. Initiating and maintaining such an import reaction is usually controllable.

A cell *performs actions* in the world like moving itself one position to another, importing from the environment substances to be used as resources in internal processes, and exporting to the environment substances that are not to be used (waste).

Table 3 External agent concepts for a cell

External agent concepts	cell
<i>Interaction with the world</i>	
observations	presence of chemicals (and their gradients) in the environment
passive observations	by passive uptake of lactose
active observations	by controlled import reactions of specific molecules from the environment; e.g., glucose
performing actions	moving by active transport of selected molecules from outside in the cell to be used as resources in internal processes exporting selected molecules to the environment as waste
<i>Communication with other agents</i>	
outgoing communication	exporting specific substances that can be picked up by another cell; external substances binding to receptors
incoming communication	importing specific substances from the environment originating from another cell

The cell *communicates to* another cell by exporting specific substances to its environment. These substances can be picked up by the other cells. The type of substance that is exported carries information about the cell's state. The cell receives *communication from* another cell by importing a specific substance that the other cell has exported to the environment, and thus gets informed about the other cell's state.

4.4 Dynamic Properties for Cell Behavior

For a more specific example, to interact with its external environment, for *E. coli* several mechanisms are known. One of these is an observation mechanism to find out which food substances are present in the environment. In the simplified case we consider we will address only glucose or lactose. Another mechanism enables *E. coli* to actively import a food substance (of its choice). This is a way of performing actions. *E. coli* makes the choice between the two types of food in the following manner.

EB1

at any point in time t
 if in trace γ at t lactose is observed, but not glucose,
 then there exists a t' with $t \leq t' \leq t+r$ such that at t' the cell will import lactose.

EB2

at any point in time t
 if in trace γ at t glucose is present,
 then there exists a t' with $t \leq t' \leq t+r$ such that at t' it will import glucose.

As before, these properties can be expressed in the formal language of TTL. However, for shortness in many cases the formal specification of the example dynamic properties are left out. From an external viewpoint (EB stands for External view on Bacterial behavior) this can be summarised and interpreted as: *E. coli* can use both substances, but it prefers glucose; it never imports both. In particular, this means that the following pattern occurs in an environment where first only lactose is present, and next glucose is added such that both are present (see trace 3 in Table 4).

Table 4 Example set of external traces of cell food import behavior

time trace	time point 0	time point 1	time point 2	time point 3	time point 4	time point 5
trace 1	lactose no glucose	lactose no glucose import lactose	lactose no glucose import lactose	no lactose no glucose import lactose	no lactose no glucose import lactose	no lactose no glucose import lactose
trace 2	no lactose glucose	no lactose glucose import glucose	lactose glucose import glucose	lactose glucose import glucose	lactose glucose import glucose	lactose glucose import glucose
trace 3	lactose no glucose	lactose no glucose import lactose	lactose glucose import lactose	lactose glucose import glucose	lactose glucose import glucose	lactose glucose import glucose
trace 4	no lactose glucose	no lactose glucose import glucose	no lactose glucose import glucose	no lactose glucose import glucose	no lactose glucose import glucose	no lactose glucose import glucose

Models in terms of dynamic properties of externally observable behavior can provide high level descriptions abstracting from what happens inside. These descriptions have their own value. However, to explain the realisation of such externally observable behavior, information is needed about what happens internally or assumptions have to be made about what happens internally. Examples of questions to be addressed in the next section then are:

- How does it actually work, how is this behavior realised; for example, how is it possible that lactose import stops while only an additional substance (glucose) occurs in the environment?
- Is there some internal decision model; how is such a decision model realised?

5 Modeling Dynamic Phenomena using Internal Agent Concepts

A description of behavior in terms of the external agent concepts abstracts from what is inside the agent. For given environmental conditions at one point in time an agent can show one behavior and at another point in time a different behavior, depending on its internal state that may differ over time. To be able to describe such dependencies of behavior on internal states, in addition to descriptions of agent states in terms of external concepts as discussed above, also descriptions of internal agent states in terms of internal concepts are useful. To obtain a high level description of such an internal agent state, the following internal agent concepts are often distinguished. Notice that these concepts provide a vocabulary for the internal processes within an agent, but still allow to abstract from the precise physical or chemical realisation of these internal states and their dynamics. Thus the concepts are generic over multiple realisations.

In Figure 2 the global structure of an agent is given with the large dotted box the boundary of the agent, outside the external agent concepts dealing with observation, action and communication, and inside the box the internal representational and motivational concepts. The latter concepts will be explained in more detail in this section. In Figure 3 a more detailed instantiation of this agent structure will be shown for the BDI-model.

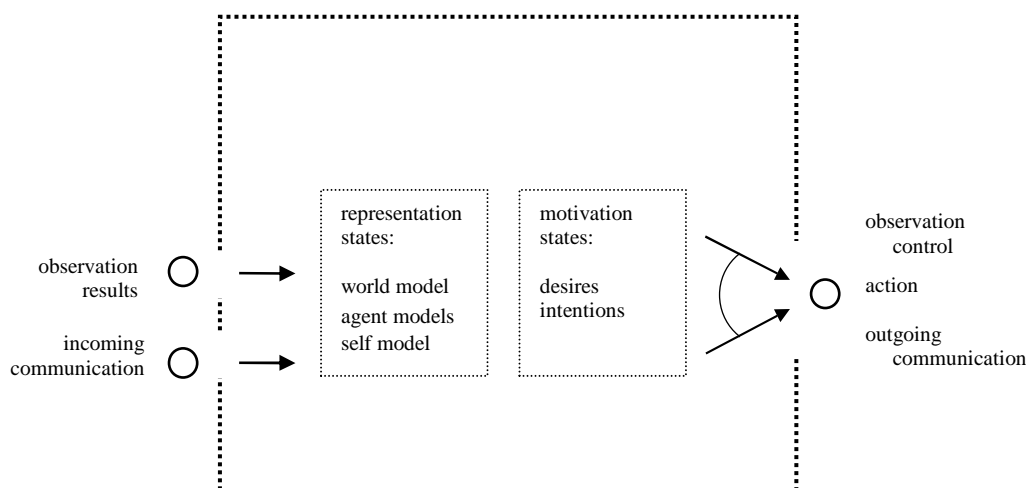


Figure 2 Global Agent Structure

5.1 Internal Agent Concepts: World Models, Agent Models and Motivations

Internal agent concepts describe the internal states and processes of an agent. In the literature the notions world model, agent model and self model are often used to describe internal states. Moreover, motivational state properties also receive much attention to describe internal states.

An agent may create and maintain information on (a model of) external *world* based on its observations of that world, on information about that world communicated by other agents, and on its (pre-existing) knowledge about the world. The agent may also create and maintain information on (models of) *other agents* in its environment based on its observations of these agents as they behave in the external world, on information about these agents communicated by other agents, and on its pre-existing knowledge about other agents. Some agents create and maintain information on (a model of) their own characteristics, internal state, and behavior.

To obtain pro-active, goal-directed behavior, an agent often represents, generates, and uses explicit motivational properties such as having a specific *desire* or *goal*, having an *intention* or *plan*. These properties guide the agent's initiatives. *Desires* are interpreted as what the agent wants to accomplish or fulfil. Agents can have desires contradictory in their fulfilment, for example desiring lots of snacks and a modest weight. Desires can be internally present, for example innate, or they can be dynamically generated, for example, on the basis of observational information. A desire, together with a sufficient additional reason (i.e., a reason in addition to the desire, which also can be viewed as a reason), leads to an intention to fulfil the desire. Such a *reason* to generate a specific intention, given a desire, is a belief in some state of affairs in the world that has to hold, in order for the intention to be generated. *Intentions* are interpreted as that the agent will make something happen (action), as soon as a belief in an *opportunity* for the action occurs, i.e., a belief in a state of affairs in the world that make successful action performance possible. These beliefs in an opportunity can be viewed as a *reason* to initiate the specific action, given the presence of the intention (i.e., a reason to initiate the action, in addition to the intention, which is also a reason). An *initiated action* can be performed *successfully* (in which case it has its expected effects) or *fail* (in which case it has no effects, or not the expected effects), depending on *enabling conditions* in the environment. Summarising, desires or goals are sources of initiative, which can be always present (static), as a kind of innate desires or goals, or dynamically generated depending on the agent's observations and history. Intentions and plans usually are dynamically generated based on desires and beliefs, to allow the agent to behave in a manner tuned to the actual circumstances in the environment.

5.2 Internal Agent Dynamics for Delayed Response Behavior

First an internal dynamics model for an animal's delayed response behavior is addressed. Next motivation-based behavior is addressed.

The following internal dynamic properties describe that an internal world model (in the form of beliefs about the world) is created and maintained within the animal, and that the animal bases its actions on this world model. The first internal dynamic property describes that the actions are based on specific internal state properties: the beliefs b and b' about the world. Here internal state property

b is the belief that food is present at position P, and internal state property b' is the belief that the screen is absent. Both can be part of the agent's world model. Based on these beliefs, the first internal dynamic property expresses that at any point in time, if both beliefs b and b' hold for the agent (i.e., b and b' are part of the agent's world model), then the agent will go to P:

IDR1

at any point in time t
 if in trace γ at t the internal state property b holds,
 and the internal state property b' holds
 then there exists a t' with $t \leq t' \leq t+r1$ such that at t' the agent will go to P

Note that within this dynamic property formulation the phrase 'the internal state property b holds' can equally well be read or written as 'the agent believes that food is present at position P', and similarly for the other belief. However, in this example the first type of phrase is used to emphasize that these beliefs are not more (and not less) than specific internal state properties, which play a mediating functional role (cf., Kim, 1996) within the internal dynamics that is usually viewed as the role of a belief.

The properties IDR2 and IDR3 describe how beliefs are created, i.e., how the agent's world model is updated in accordance to changes in the world. For example, property IDR2 expresses how a specific interaction with the environment (the observation that food is present at a specific position P) leads to a specific internal state property (the belief b) as part of the agent's world model.

IDR2

at any point in time t
 if in trace γ at t the agent observes that food is present at position P,
 then there exists a t' with $t \leq t' \leq t+r2$ such that at t' internal state property b will hold

IDR3

at any point in time t
 if in trace γ at t the agent observes the absence of the screen,
 then there exists a t' with $t \leq t' \leq t+r3$ such that at t' internal state property b' will hold

Finally, internal dynamic properties IDR4 and IDR5 express that beliefs (as internal state properties, part of the agent's world model) will persist over time as long as no opposite observations are received. Here IDR4 expresses that after belief b started to hold, it will persist until the opposite observation is made, and similarly for belief b' in IDR5.

IDR4

for all time points t1 and t2 with $t1 < t2$
 if in trace γ at t1 internal state property b holds,
 and between t1 and t2 the agent does not observe that food is not present at position P,
 then at t2 internal state property b holds

IDR5

for all time points t_1 and t_2 with $t_1 < t_2$

if in trace γ at t_1 internal state property b' holds,

and between t_1 and t_2 the agent does not observe that food is not present at position P,

then at t_2 internal state property b' holds

5.3 Internal Agent Dynamics for Motivation-Based Behavior

An example of a motivation-based model for internal dynamics of the animal is also based on the internal state properties b , b' (part of the agent's world model) as above, but in addition on the internal state properties d and i (d for desire for food and i for intention to go to P; these internal state properties are not part of the agent's world model). The following dynamics are assumed (see Figure 3).




Figure 3 Internal dynamics in a BDI-model for motivation-based behavior

Here the first two internal dynamic properties express under which circumstances a desire leads to an intention (IMB1), and under which circumstances an intention leads to an action (IMB2). Here IMB stands for Internal view on Motivation-Based Behavior.

IMB1

at any point in time t

if in trace γ at t the internal state property d holds,

and the internal state property b holds

then there exists a t' with $t \leq t' \leq t+r$ such that at t' the internal state property i will hold

IMB2

at any point in time t

if in trace γ at t the internal state property i holds,

and the internal state property b' holds

then there exists a t' with $t \leq t' \leq t+r$ such that at t' the agent will go to P

Note that in these dynamic properties the phrase ‘the internal property d holds’ can equally well be replaced by ‘the agent desires food’, and ‘the internal property i holds’ can equally well be replaced by ‘the agent intends to go to P’.

Further dynamic properties that express how observations lead to beliefs (IDR2 and IDR3) and how beliefs persist over time (IDR4 and IDR5) can be taken from Section 5.2 above. Persistence of the intention can be modelled in different manners. For example, an intention may be required to persist even if the desire already disappeared: single-minded; see, for example, (Brazier et al., 1999). Another possibility is that the intention is required to persist as long as the desire persists, but may disappear as soon as the desire has disappeared: open minded; see (Brazier et al., 1999). This latter possibility is expressed in dynamic property IMB3.

IMB3

for all time points t_1 and t_2 with $t_1 < t_2$
 if in trace γ at t_1 internal state property i holds,
 and between t_1 and t_2 internal state property d holds,
 then at t_2 internal state property i holds

This description of the animal behavior is much more abstract and transparent than a description in terms of the physiological processes, which in general is intractable. However, a relevant question is whether this model has a foundation. It can be grounded either in its relationship to the externally observable behavior (*instrumentalist view*) or in its relationship to internal physiological or biochemical (e.g., brain) processes (*realist view*). The foundational issue how for an instrumentalist view the properties describing internal dynamics are to relate to the externally observable behavior will be discussed in more detail in Section 6. The realist view will be discussed in more detail in Section 5.5 in the context of cellular processes.

5.4 Internal Agent Dynamics for Relative Adaptive Behavior

Another example of a model for internal dynamics can be obtained for the relative adaptivity property (RA). This dynamic property can be realised by different specific sets of dynamic properties. For example, in the case of a discrete time frame, some internal state property (skill capacity c) can be assumed which persists but is subject to a certain decay (during time periods that no training is performed), indicated by parameter r (with value between 0 and 1). One set of such internal dynamic properties is the following:

GS1

At any point in time t ,
 if in trace γ at t the exercising intensity is v , and the skill capacity has value c ,
 then at time $t + 1$ the skill capacity will have value $r.c + (1-r).v$

In formalised form:

$$\forall t \ [\text{state}(\gamma, t, \text{InOnt}(A)) \models \text{has_intensity}(\text{exercising}, v) \ \& \ \text{state}(\gamma, t, \text{InOnt}(A)) \models \text{has_value}(\text{skill_capacity}, c) \\ \Rightarrow \ \text{state}(\gamma, t+1, \text{InOnt}(A)) \models \text{has_value}(\text{skill_capacity}, r.c + (1 - r).v)]$$

This internal dynamic property specifies how the internal state property indicating the agent's skill capacity is updated based on a specific exercising intensity, and on decay over time.

GS2

At any point in time t ,
 if in trace γ at t the skill capacity has value c ,
 then at time $t + 1$ the demonstrated skill level is $a.c$

In formalised form:

$$\forall t \ [\text{state}(\gamma, t, \text{InOnt}(A)) \models \text{has_value}(\text{skill_capacity}, c) \Rightarrow \ \text{state}(\gamma, t+1, \text{OutOnt}(A)) \models \text{has_level}(\text{demo_skill}, a.c)]$$

This internal dynamic property specifies how a certain skill capacity leads to a certain demonstrated skill level. Here a is a parameter (with positive value) that indicates how (in which proportion) capacity affects skill.

5.5 Intracellular Dynamics in Terms of Internal Agent Concepts

A cell as *E. coli* behaves in a manner that differs for different environmental circumstances. Therefore, to obtain a description of a cell from an agent perspective, questions such as the following are to be addressed:

- How does *E. coli* take the environment into account in its decisions; does it somehow know about the environment (world model); is some sensing mechanism available to build and maintain such a world model?
- How does it make the appropriate choices in behavior; how can motivational properties play a role; is a model based on beliefs, desires and intentions applicable?

First, the notions world model, agent model and self model are addressed, next motivational properties.

Table 5 Internal agent concepts for a cell

Internal agent concepts	<i>cell</i>
<i>World Model</i>	presence of internal chemical concentrations relating to the current state of the external environment
<i>Agent Models</i>	presence of internal chemical concentrations relating to the external presence of other cells and their specific states
<i>Self Model</i>	presence of genetic information in DNA
<i>Desires</i>	statically: built-in genes in DNA, e.g., to grow and to divide dynamically: activated (depending on circumstances) genes in the DNA
<i>Intentions</i>	presence of specific mRNA and enzymes playing a role in the control of specific intracellular processes

The concept ‘world model’ means that internal state properties exist that can be considered as representations (in some form) of properties of the state of the external world. Put differently, certain internal state properties can be granted ‘representational content’ in the sense that their occurrence indicates that a certain state of affairs is present in the external world. Under normal conditions these representations are assumed to be sincere with respect to the actual properties of the external world state. This means that there is a clear correlation between the internal state property that is the representation and the external world state property. The causal/correlational approach assumes that, to realize such a correlation, causal chains exist between the external world state property and the internal representation, e.g., (Kim, 1996), pp. 191-193; (Jacob, 1997), pp. 94-104. Within cell biology, indeed such causal chains are known as chemical pathways from the environment to within the cell. For example, such causal chains justify to identify representational content of the presence of an internal concentration of CRPcAMP above a certain level as ‘glucose being absent in the external environment’, and of the internal presence of a certain concentration of lactose as ‘lactose being present in the external environment’. This shows ways in which a cell is able to build and maintain a *world model*.

The above accounts for ways how the cell can have internal images of the relevant aspects of the external environment, enabling the cell to control its intracellular processes in an informed manner. Similar mechanisms exist to obtain internal representations of the fact that other cells in a certain state are present in the environment, given that these cells export substances to their environment that indicate their presence and in what state they are. If these substances are picked up by a cell, then based on the causal/correlational approach, in a similar way as described above, an internal model of other cells (*agent models*) are created. Using these agent models, the cell is able to control its intracellular processes in a manner that is also informed on the states of other cells in its environment. In its genetic material, the cell has an explicit representation of its own characteristics; this can be considered as having a *self model*. Via transcriptional and translational processes, this self model has a major influence on the intracellular processes.

As examples of *desires* of a bacterium, built-in (in the DNA) desires to grow and to divide can be considered. However, the desires of a cell can also get more specific forms if the cell specialises (differentiation) in a form that makes it appropriate for a specific task. In such a case the desire can be considered as built-in in the differentiated structure of the cell. Yet another way of considering desires in a cell is in a dynamic manner, when on the basis of information on its environment certain genes and transcriptional processes become active.

Intentions can be considered present in the cell in that, given its desires and depending on the observed environment it is able to make an informed choice between alternatives of specific (packages of) processes; e.g., import to provide resources for a specific type of metabolism, followed by the type of metabolic process. Such a package of tuned processes can be viewed as a built-in intention or plan alternative; usually a number of them are available. In an informed and dynamic manner, also involving its desires, the cell chooses between the alternatives.

As a more detailed example, consider the following internal state properties for *E. coli* involved in its behavior to import lactose.

- the belief that no glucose is present in the environment
- the belief that lactose is present in the environment
- the desire to grow
- the desire to import lactose
- the intention to import lactose

If the following relationships between these internal state properties are assumed (these are kept simple to explain the idea), for an environment in which lactose but not glucose is present, the lactose import behavior can be explained. First the generation of desires.

IBB0 Innate desire

at any point in time,
the bacterium desires to grow

IBB1 To more specific desires

at any point in time t ,
if in trace γ at t the bacterium desires to grow,
then there exists a t' with $t \leq t' \leq t+r1$ such that at t' it will desire to import lactose

Note that in these dynamic properties a phrase such as ‘the bacterium desires to grow’ could equally well be replaced by a phrase such as ‘the internal state property desire(grow) holds’. In Section 5.5 the latter type of phrases were used; in the current example the former type of phrases are used for readability. Next, the generation of an intention based on a desire and an (additional) reason.

IBB2 Desire and (additional) reason together lead to intention

at any point in time t ,
if in trace γ at t the bacterium desires to import lactose,
and it believes that lactose is externally present,
and it believes that glucose is not externally present,
then there exists a t' with $t \leq t' \leq t+r2$ such that at t' it will intend to import lactose

The generation of an action based on an intention and a belief in an opportunity.

IBB3 Intention and belief in opportunity together lead to action initiation

at any point in time t
if in trace γ at t the bacterium intends to perform import of lactose
and it believes that lactose is externally present,
and it believes that no glucose is externally present
then there exists a t' with $t \leq t' \leq t+r3$ such that at t' it will initiate import of lactose

Notice that the beliefs in the opportunity to perform the action are the same as the beliefs that make up the reason (in addition to the desire) to generate the intention (property **IBB2**), i.e., the belief that there is lactose but not glucose in the environment. This is not necessarily redundant, because the intention may be generated and persist over a longer time period, whereas the environment may have changed in the meantime; the conditions on these beliefs may have a different outcome for a different point in time. The initiated action only succeeds if at execution time lactose is externally present (enabling condition). If this enabling condition is not fulfilled, the expected effect (lactose inside) will not occur.

IBB4 Action initiation and enabling condition together lead to action effects

at any point in time t
 if in trace γ at t the bacterium initiates import of lactose
 and lactose is externally present,
 then there exists a t' with $t \leq t' \leq t+r4$ such that at t' lactose will be present internally

In addition to the above, belief creation or sensing properties are needed:

IBB5 Sensing presence of lactose

at any point in time t ,
 if in trace γ at t the bacterium has lactose in its environment,
 then there exists a t' with $t \leq t' \leq t+r5$ such that at t' it believes that lactose is externally present

Notice that this internal dynamic property only specifies that at least one time point exists at which the cell believes that lactose is externally present. To specify that this belief persists over longer time durations, dynamic persistence properties for the beliefs, similar to **IDR4** and **IDR5** in Section 5.2, are assumed.

IBB6 Sensing absence of lactose

at any point in time t ,
 if in trace γ at t the bacterium has no lactose in its environment,
 then there exists a t' with $t \leq t' \leq t+r6$ such that at t' it believes that lactose is not externally present

IBB7 Sensing presence of glucose

at any point in time t ,
 if in trace γ at t the bacterium has glucose in its environment,
 then there exists a t' with $t \leq t' \leq t+r7$ such that at t' it believes that glucose is externally present

IBB8 Sensing absence of glucose

at any point in time t ,
 if in trace γ at t the bacterium has no glucose in its environment,
 then there exists a t' with $t \leq t' \leq t+r8$ such that at t' it believes that glucose is not externally present

Note that for reasons of presentation, here these relationships are kept simple. They can be made more sophisticated, for example, by involving time durations for antecedent and consequent, and more restricted response time intervals. Moreover, also internal state properties involving the import of other substances, such as glucose, can be added.

The different maximal response times r_1 to r_8 within the model leave some freedom. Viewed from an instrumentalist perspective, a constraint is that the overall cumulative delay implied by these internal delay parameters should be realistic compared to the response time shown in externally observable behavior. Viewed from a realist perspective more detailed information can be gained. For the realist perspective, the aim is to find physical or chemical counterparts of the beliefs, desires, intentions, and actions, and show that they are (causally) related in a manner justifying assumed relationships such as $\text{IBB}_0, \dots, \text{IBB}_8$ (and other relationships involving, for example, glucose import). For a more elaborate version of this case study, this actually has been done; see (Jonker, Snoep, Treur, Westerhoff and Wijngaards, 2002). For example, the desire to import lactose was associated to ‘lactose import operon internally present in DNA concentration above 0.01 mmol/l’, the intention to perform lactose import has been associated to ‘lactose import mRNA internally present, concentration above 0.1 mmol/l’, and the action to perform lactose import has been associated to ‘active lactose import enzymes internally present, concentration above 0.1 mmol/l’.

From these one-to-one correspondences between the beliefs, desires and intentions and actions on the one hand, and chemical conditions within the cell on the other hand, more can be said about the response times r_1 to r_8 above. For example, the dynamic property IBB_2 with its maximal response parameter r_2 relates to the process of transcription to obtain a relevant amount of mRNA, which may take 60 seconds, whereas the generation of a relevant amount of an active enzyme specified by this mRNA (relating to property IBB_3 with its maximal response parameter r_3) may also take, e.g., 20 seconds. Also, realistic values for r_4 to r_8 are within the 0.1 second range. Moreover, a realistic value for the response time r_1 within property IBB_1 is 0, since the presence of DNA directly implies the presence of the lactose operon on this DNA. These differences in response times provide additional insight in the dynamics of the model as described; this shows an advantage of a realist approach. For example, given the relatively slow response of property IBB_2 , it is certainly not redundant to have a renewed check on the beliefs in property IBB_3 .

6 Relationships between an Agent’s Internal Dynamics and its Behavior

In Sections 4 and 5 agent-oriented models have been discussed for the externally observable agent behavior, and for the internal agent processes, respectively. In this section the relationships between the two types of models is addressed.

6.1 Relationships between Internal Dynamics and Externally Observable Behavior

Two agent-oriented views on dynamics of an agent are possible:

- a view VE restricted to externally observable behavior: based on InteractionOnt
- a view VI including the internal dynamics: based on IntOnt \cup InteractionOnt

VE and VI each define a set of traces which are restrictions of the overall traces from Traces(OvOnt) :

$$\begin{aligned} \text{Traces(VE)} &= \{ \gamma | \text{InteractionOnt} \mid \gamma \in \text{Traces(OvOnt)} \} \\ \text{Traces(VI)} &= \{ \gamma | \text{IntOnt} \cup \text{InteractionOnt} \mid \gamma \in \text{Traces(OvOnt)} \} \end{aligned}$$

The two views can be related in the following manner. If a given trace is in accordance with the view of the internal dynamics, then its externally observable behavior should be in accordance with the externally observable view. In other words, if an overall trace γ of a process is given, and γ is in accordance with VI, then the restriction of γ to the interaction ontology should be in accordance with VE, or:

$$\{ \gamma | \text{InteractionOnt} \mid \gamma \in \text{Traces(VI)} \} \subseteq \text{Traces(VE)}$$

Suppose for each of these views characterising dynamic properties have been identified: the sets DPE and DPI for external, resp. internal properties. According to the above relationship it may be reasonable to expect that if a trace γ for IntOnt \cup InteractionOnt satisfies the set of internal dynamic properties DPI (i.e., the agent's internal processes function according to DPI), this implies that γ satisfies the set of externally observable dynamic properties DPE (i.e., the agent behaves according to DPE). If the properties of DPI are enumerated by l_1, \dots, l_m , and those of DPE by E_1, \dots, E_n , then such a relationship can be expressed as follows:

$$l_1 \& \dots \& l_m \Rightarrow E_j \text{ for all } j$$

Within this (logical) relationship it may be the case that for a specific dynamic property E_j only a subset of the set of properties DPI is needed. In that case the antecedent can contain only the relevant elements of DPI. Such logical relationships between dynamic properties can be exploited for mathematical (i.e., formal logical) analysis. Moreover, explanations, based on the internal properties, of dynamic properties of the externally observable behavior become justified: these dynamic properties of externally observable behavior can be *deduced* from internal dynamic properties used in such an explanation (one of the criteria for Deductive-Nomological or D-N explanations in (Hempel and Oppenheim, 1948)). This general pattern will be discussed in more detail for each of the two case studies developed in previous sections.

6.2 Relationships between Internal and External View for Delayed Response Behavior

The relationships between internal and external dynamic properties are discussed for response behavior and for relative adaptive behavior in the animal example.

Delayed response behavior as described in Section 4.2 can be explained from its internal dynamics by assuming that the animal creates and maintains a world model as a form of memory, as described in Section 5.2. In particular, assuming $r \geq r_1 + \max(r_2, r_3)$, the following relationship between the dynamic properties holds:

$$\text{IDR1} \& \text{IDR2} \& \text{IDR3} \& \text{IDR4} \& \text{IDR5} \Rightarrow \text{EDR2}$$

Such relationships are useful for formal analysis. For example, in this case an explanation of externally observable behavior can be made as follows. Suppose particular environmental information B is given that food has been visible, and the screen has been taken away, and that under these environmental circumstances the external properties entail that the animal goes to P.

Why does the animal go to P?

The animal goes to P because it *believed* that food is present at P, and it *believed* that no screen is present.

(IDR1)

Why does it believe that food is present and that no screen is present?

It had the *belief* that food is present because it *observed* that food is present, and after that it did not *observe* that no food is present. (IDR2, IDR4)

It had the *belief* that no screen is present because it *observed* that no screen is present, and after that it did not *observe* that a screen is present. (IDR3, IDR5)

As indicated, each step of this explanation relates to one of the internal dynamic properties. Taking the internal properties as laws, this fits in the Deductive-Nomological (D-N) model for explanation (Hempel and Oppenheim, 1948): from the environmental circumstances B and IDR1 & IDR2 & IDR3 & IDR4 & IDR5 it can be *deduced* that the animal goes to P. To be more specific, this is an example of an instrumentalist explanation, and fits also in what Dennett (1987, 1991) calls the *intentional stance*, which was put forward as opposed to explanation from a direct physical perspective (the *physical stance*). According to Dennett (1987) an *intentional system* is an entity whose behavior can be predicted by the method of attributing beliefs, designs and rational acumen. Intentional notions such as *beliefs, desires, intentions, commitments, goals, plans, preference* may be assigned to such systems.

An interesting philosophical discussion has taken place on what the justification of an intentional stance explanation (including the underlying assumptions) such as the one above is. In particular, the status of the internal state properties as used can be a point of debate. As already mentioned, two different positions taken are the *instrumentalist position* and the *realist position*. Dennett is considered closest to the instrumentalist position. This position claims that the only source of justification that is necessary is the success of predictions and explanations. There is no need to assume or to show that the internal properties have some counterpart in the physical (or chemical) world. They can be considered intermediate steps; the only thing that counts is that the predictions or explanations have correct outcomes. In contrast, the realist position claims that the intermediate state properties themselves also need to be related to and follow physical reality in a correct manner.

The advantage of intentional stance (or folk-psychological) explanations, is explained by Dennett, (1991) as follows :

'Predicting that someone will duck if you throw a brick at him is easy from the folk-psychological stance; it is and will always be intractable if you have to trace the protons from brick to eyeball, the neurotransmitters from optic nerve to motor nerve, and so forth.' (Dennett, 1991), p. 42

not [GS1 & GS2 \Rightarrow IDS]

For example, if at some point in time the exercising intensity is 0, and at that time point the demonstrated skill level in γ_2 is strictly higher than the demonstrated skill level in γ_1 , then assuming property GS1 with $0 < c < 1$, the difference will become less, which contradicts property IDS.

6.4 Relationships between Internal and External View for the Cell

How can the dynamic properties on internal processes be used to explain externally observable behavior? Assuming the internal state properties and their relations IBB0, ..., IBB8 as described above, an example of such an explanation – in iterated form – is as follows:

Why does this *E. coli* import lactose?

This *E. coli* imports lactose because it had the *intention* to perform lactose import and it *believed* that no glucose is present in the environment.

Why the intention to perform lactose import and why that belief?

It had the *intention* to perform lactose import because it had the *desire* to import lactose, and it *believed* that lactose is present in the environment and that glucose is not present.

It had the *belief* that no glucose is present externally because it *observed* that no glucose was present.

Why the desire to import lactose and why the belief that lactose was present ?

It had the *desire* to import lactose because it has the *desire* to grow.

It had the *belief* that lactose is present because the bacterium *observed* that lactose is present.

Why the desire to grow?

This *desire* is innate (in the bacterium's DNA).

As for the case study of animal behavior, this type of description fits in the *intentional stance*. From an instrumentalist perspective, a justification can be found in the same way as in Section 6.2. To define more precisely what the criterion is that should be fulfilled, the dynamic properties EB1 and EB2 from Section 4.3 can be used. It can be shown that all explanations or predictions that can be generated (under all different appropriate possible environmental conditions and histories), using relationships such as IBB0, ..., IBB8 (augmented with other relationships involving, for example, glucose import), indeed entail such a specification, as long as the response time r in EB1 is taken sufficiently large. More specifically, assuming $r \geq r_1 + r_2 + r_3 + r_4 + \max(r_5, r_6, r_7, r_8)$, the following logical relationship holds

IBB0 & IBB1 & IBB2 & IBB3 & IBB4 & IBB5 & IBB6 & IBB8 \Rightarrow EB1

Establishing such a logical relationship (which can be done by mathematical proof, or even by an automated theorem prover or checker) justifies the explanation given above from an instrumentalist perspective.

In the *E. coli* case study also a *realist* justification of the explanation is possible: based on direct relationships between on the one hand the internal state properties beliefs, desires, and intentions in the explanation and on the other hand certain (physical or chemical) internal state properties of this

cell, as discussed in Section 5.5. If the chemical counterparts of the state properties fulfil dependence relations covering the dynamic properties $\text{IBB}_0, \dots, \text{IBB}_8$, then this provides a direct justification of each of these properties. So, in general, from the realist perspective the aim is to find physical or chemical counterparts of the beliefs, desires and intentions, and show that they are (causally) related in a manner justifying assumed relationships such as $\text{IBB}_0, \dots, \text{IBB}_8$ (and other relationships involving, for example, glucose import). For the example, this actually has been done; see (Jonker, Snoep, Treur, Westerhoff and Wijngaards, 2002); see also Section 5.5. With these (and some other) associations, the above relationships are covered by biochemical principles, thus obtaining a basis for realist explanations.

7 Discussion

In this section some reflection is made on the level of conceptualisation and explanation of phenomena. Many authors, for example from the area of Philosophy of Mind, feel that the more complex the phenomena, the less adequate are explanations directly at a physical (or chemical) level; higher level concepts are needed to manage the complexity. In this discussion first a discussion is presented on some positions from Philosophy. Next it is discussed how concepts and techniques from the agent system area provide high level descriptions of dynamics.

7.1 Agent-Oriented Modeling and Complexity of Biological Processes

The area of agent systems provides higher level concepts and description languages to model complex phenomena, and in particular the dynamic or behavioral aspects of such phenomena. The concepts and modeling techniques for *externally observable behavior* introduced in Section 4 abstract from the specific physical states and dynamics exemplifying the behavior, and the specific mechanisms behind the internal processes realising the observable behavior. Thus a high level modeling approach is obtained, providing descriptions of the essence of externally observable behavioral patterns in terms of dynamic properties that are abstract and generic with respect to multiple physical or chemical occurrences of essentially the same behavioral pattern. Thus, a species-independent categorisation of specifications of different types of behavior can be obtained, such as

- stimulus-response behavior,
- delayed response behavior,
- goal-directed or motivation-based behavior,
- adaptive behavior.

In particular for comparative studies of behavior such species-independent specifications are useful.

From an agent oriented perspective also high level models for *internal processes within an agent* are contributed. These models introduced in Section 5 provide details about internal processes leading to externally observable behavior, but still abstract from the physical and chemical details of the mechanisms behind the internal processes. Thus also for the essence of internal processes

abstract generic descriptions in terms of dynamic properties are offered that are multiply realisable in the internal physics and chemistry.

Both type of agent-oriented descriptions, for the dynamics of internal processes and for externally observable behavior can be related to each other. In particular dynamic properties characterising the internal dynamics, together imply dynamic properties characterising the externally observable behavior. Such logical relationships provide a solid basis for explanation of externally observable behavior by internal dynamics.

In comparison, for example in the area of modeling intracellular processes, the most widely used approach is based on differential equations, which are integrated numerically (Westerhoff, 2001). For some small unicellular organisms, a few isolated chemical pathways are understood in sufficient kinetic detail to obtain a description of their import and primary processing of nutrients; e.g., in *Escherichia coli* (Rohwer et al., 2000, Wang *et al.*, 2001, Ben-Jacob *et al.*, 1997), or yeast (Teusink *et al.*, 2000; Rizzi et al., 1997). However, this approach is felt to have serious limitations in tackling more large-scale cellular systems. First, hundreds or more reaction parameters are needed, for which reliable values are rarely available (Teusink *et al.*, 2000; Kholodenko *et al.*, 1999). This can seriously compromise the feasibility of the general model. Second, actual behavior of intracellular pathways may be much less complex than is theoretically possible on the basis of the complexity of the chemical processes (e.g., Rotterdam *et al.*, 2002). For example, the decision model for food import of *E. coli* discussed in Sections 4.3 and 5.5 has a much lower complexity than the differential equation model for all biochemical reactions involved in these decisions.

At best, and only if all system parameters and internal connections are known and sufficiently tuned, the traditional approach delivers a computer replica of (part of) the living cell, which however is almost as remote from human understanding as the target system itself. This is because the modeling approach requires a description that is complete, inherently low-level, detailed and complex. In contrast, the human mind operates by abstraction in order to understand an essence.

To get a numerical impression, in a (modest) biochemical model of a cell taking into account 100 substances, for which concentrations are assumed to be measured in real numbers, say with an accuracy of 3 digits, the size of the space of all possible internal states is about $1000^{100} = 10^{300}$. In comparison, in a motivation-based model of the same cell consisting of 30 internal state properties (beliefs, desires, intentions), the size of the space of all possible internal states is about $2^{30} = 10^9$. This state space is a multitude of orders of magnitude smaller. This also has a drastic impact on the number of possible trajectories through such a state space.

7.2 Support for Agent-Oriented Analysis and Simulation

In the case studies the emphasis of the modeling has been on *analysis*. The analysis method for the dynamics may involve the following ingredients:

- Specify state properties and dynamic properties of the overall process
- Identify the agents within the overall process
- Specify state properties and dynamic properties for the behavior of these agents

- Relate the dynamic properties of the overall process to the dynamic properties for agent behavior (using some organisational model)
- Specify state properties and dynamic properties for the internal processes within agents
- Relate dynamic properties for agent behavior to dynamic properties of the internal agent processes
- Relate state properties to physical or chemical state properties
- Relate dynamic properties to physical or chemical dynamic properties
- Check given traces of dynamics against dynamic properties

Software support for some of these items within analysis has been developed or is under development. For example, an editor to specify dynamic properties according to a specific format, a theorem prover that relates (specific types of) dynamic properties to other dynamic properties, and a (model) checker that checks whether dynamic properties hold in a given trace; e.g., (Jonker, Letia, and Treur, 2002; Jonker, Treur, and Wijngaards, 2003).

The agent-oriented modeling perspective presented in this paper, can also be used as a basis for *simulation*. The idea, which is an extension of the paradigm of executable temporal logic (Barringer et al., 1996; Fisher 1996, 2005) is as follows. Specify basic *executable properties* in a format

$$\text{past \& current state} \Rightarrow \text{next or future state}$$

Actually, most if not all of the dynamic properties encountered in this paper can be written in this format. Then a trace can be developed by starting at time $t = 0$ and for each time point up to which the trace already has been constructed, checking which antecedents of executable properties hold in the already constructed trace. For these executable properties, add the consequent to the trace, i.e., extend the trace in time in such a manner that the consequent holds. Also for simulation based on executable dynamic properties a supporting software environment has been developed; e.g., (Jonker, Treur, and Wijngaards, 2003).

Acknowledgements

The authors have learned a lot of this area from discussions and cooperation with Jaap Heringa, Jacky Snoep, Hans Westerhoff, Wouter Wijngaards. The text was further improved on the basis of useful comments of anonymous referees.

References

- Barringer, H., M. Fisher, D. Gabbay, R. Owens, and M. Reynolds (1996). *The Imperative Future: Principles of Executable Temporal Logic*, Research Studies Press Ltd. and John Wiley & Sons.
- Ben-Jacob, E., Cohen, I., Czirók, A., Vicsek, T., and Gutnick, D.L. (1997). Chemomodulation of Cellular Movement and Collective Formation of Vortices by Swarming Bacteria. In: *Physica A*, Vol. 238, pp. 181-197.
- Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J., (2005). LEADSTO: a Language and Environment for Analysis of Dynamics by SimulaTiOn. In: Eymann, T., Kluegl, F., Lamersdorf,

- W., Klusch, M., and Huhns, M.N. (eds.), *Proc. of the Third German Conference on Multi-Agent System Technologies, MATES'05*. Lecture Notes in Artificial Intelligence, vol. 3550. Springer Verlag, 2005, pp. 165-178.
- Bosse, T., Jonker, C.M., and Treur, J., (2006). On the Use of Multi-Agent Organisation Modelling Techniques to Address Biological Organisation. Technical Report. Vrije Universiteit Amsterdam.
- Brazier, F.M.T., Dunin-Keplicz, B., Treur, J., and Verbrugge, L.C. (1999). Modelling Internal Dynamic Behavior of BDI agents. In: J.-J. Ch. Meyer and P.Y. Schobbies (eds.), *Formal Models of Agents (Selected papers from final ModelAge Workshop)*. Lecture Notes in AI, vol. 1760, Springer Verlag, pp. 36-56.
- Brazier, F.M.T., Jonker, C.M., and Treur, J., (2002). Principles of Component-Based Design of Intelligent Agents. *Data and Knowledge Engineering*, vol. 41, 2002, pp. 1-28.
- Brazier, F.M.T., Jonker, C.M., and Treur, J. (2000). Compositional Design and Reuse of a Generic Agent Model. *Applied Artificial Intelligence Journal*, vol. 14, 2000, pp. 491-538.
- Dardenne, A., Lamsweerde, A. van, and Fickas, S. (1993). Goal-directed Requirements Acquisition. *Science in Computer Programming*, vol. 20, pp. 3-50.
- Dennett, D.C. (1991). Real Patterns. *The Journal of Philosophy*, vol. 88, 1991, pp. 27-51.
- Dennett, D. (1987). *The Intentional Stance*, MIT Press, Cambridge, MA.
- Dubois, E., Du Bois, P., and Zeippen, J.M. (1995). A Formal Requirements Engineering Method for Real-Time, Concurrent, and Distributed Systems. In: Proceedings of the Real-Time Systems Conference, RTS'95.
- Fisher, M. (2005). Temporal Development Methods for Agent-Based Systems, *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 10, 2005, pp. 41-66.
- Fisher, M. (1996). A Temporal Semantics for Concurrent METATEM, *Journal of Symbolic Computation* 22(5):627-648, November/December 1996, Academic Press.
- Galton, A. (2003). Temporal Logic. *Stanford Encyclopedia of Philosophy*,
URL: <http://plato.stanford.edu/entries/logic-temporal/#2>.
- Galton, A. (2006). Operators vs Arguments: The Ins and Outs of Reification. *Synthese*, vol. 150, 2006, pp. 415-441.
- Gamboa, R. (2000). Continuity and Differentiability in ACL2. In M. Kaufmann, P. Manolios, and J.S. Moore (eds.), *Computer-Aided Reasoning: ACL2 Case Studies*. Kluwer Academic Press.
- Gamboa, R., and Kaufmann, M. (2001). Nonstandard Analysis in ACL2. *Journal of Automated Reasoning*, vol. 27, pp. 323-351.
- Georgeff, M.P., and Lansky, A.L. (1987). Reactive Reasoning and Planning. Proc. of the National Conference of the American Association for AI, AAAI'87. Morgan Kaufman.
- Hempel, C.G. and Oppenheim, P. (1948), Studies in the Logic of Explanation. *Philosophy of Science*, vol. 15, pp. 135-175.
- Herlea Damian, D.E., Jonker, C.M., Treur, J., and Wijngaards, N.J.E., (2005). Integration of Behavioural Requirements Specification within Compositional Knowledge Engineering. *Knowledge-Based Systems Journal*, vol. 18, 2005, pp. 353-365.

- Herlea, D.E., Jonker, C.M., Treur, J., and Wijngaards, N.J.E. (1999). Specification of Behavioral Requirements within Compositional Multi-Agent System Design. In: F.J. Garijo, M. Boman (eds.), *Multi-Agent System Engineering, Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'99*. Lecture Notes in AI, vol. 1647, Springer Verlag, Berlin, 1999, pp. 8-27.
- Jacob, P. (1997). *What Minds Can Do: Intentionality in a Non-Intentional World*. Cambridge University Press.
- Jennings, N.R. (2000). On Agent-Based Software Engineering. *Artificial Intelligence*, vol. 117, pp. 277-296.
- Jennings, N.R., and Wooldridge, M. (1998a), Applications of Intelligent Agents. In: (Jennings and Wooldridge, 1998b), pp. 3-28.
- Jennings, N.R., and Wooldridge, M. (eds.) (1998b), *Agent Technology: Foundations, Applications, and Markets*. Springer Verlag.
- Jonker, C.M., Letia, I.A., and Treur, J., (2002). Diagnosis of the Dynamics within an Organisation by Trace Checking of Behavioural Requirements. In: Wooldridge, M., Weiss, G., and Ciancarini, P. (eds.), *Proc. of the 2nd International Workshop on Agent-Oriented Software Engineering, AOSE'01*. Lecture Notes in Computer Science, vol. 2222. Springer Verlag, 2002, pp. 17-32.
- Jonker, C.M., Snoep, J.L., Treur, J., Westerhoff, H.V., and Wijngaards, W.C.A. (2001). Temporal Modelling of the Dynamics of Intentional Bacterial Behaviour. In: R. Cottam and R. Paton (eds.), *Proc. of the Fourth International Workshop on Information Processing in Cells and Tissues, IPCAT 2001*, 2001.
- Jonker, C.M., Snoep, J.L., Treur, J., Westerhoff, H.V., and Wijngaards, W.C.A. (2002). Putting Intentions into Cell Biochemistry: An Artificial Intelligence Perspective. *Journal of Theoretical Biology*, vol. 214, pp. 105-134.
- Jonker, C.M., and Treur, J., (2002). Compositional Verification of Multi-Agent Systems: a Formal Analysis of Pro-activeness and Reactiveness. *International Journal of Cooperative Information Systems*, vol. 11, 2002, pp. 51-92. Earlier, shorter version in: W.P. de Roeper, H. Langmaack, A. Pnueli (eds.), *Proceedings of the International Workshop on Compositionality, COMPOS'97*. Lecture Notes in Computer Science, vol. 1536, Springer Verlag, 1998, pp. 350-380.
- Jonker, C.M., and Treur, J. (1998b), Agent-based Simulation of Reactive, Pro-active and Social Animal Behavior. In: J. Mira, A.P. del Pobil, and M. Ali (eds.), *Methodology and Tools in Knowledge-Based Systems (Proceedings of the 11th International Conference on Industrial and Engineering Applications of AI and Expert Systems, IEA/AIE'98, vol. I)*, Lecture Notes in AI, vol. 1415, Springer Verlag, 1998, pp. 584-595.
- Jonker, C.M., and Treur, J., (2003). A Temporal-Interactivist Perspective on the Dynamics of Mental States. *Cognitive Systems Research Journal*, vol. 4, 2003, pp. 137-155.
- Jonker, C.M., Treur, J., and Wijngaards, W.C.A., (2003). A Temporal Modelling Environment for Internally Grounded Beliefs, Desires and Intentions. *Cognitive Systems Research Journal*, vol. 4(3), 2003, pp. 191-210. Earlier, shorter version in: B. Dunin-Keplicz and E. Nawarecki (eds.), *From Theory to Practice in Multi-Agent Systems, Proc. of the Second International Workshop of*

- Central and Eastern Europe on Multi-Agent Systems, CEEMAS'01*, 2001. Lecture Notes in AI, vol. 2296, Springer Verlag, 2002.
- Kholodenko, B.N., Demin, O.V., Moehren, G., Hoek, J.B. (1999). Quantification of short term signaling by the epidermal growth factor receptor. *J. Biol Chem.*, 274(42):30169-81.
- Kim, J. (1996). *Philosophy of Mind*. Westview press.
- Müller, J.P. (1996). *The Design of Intelligent Agents: a Layered Approach*. Lecture Notes in AI, vol. 1177, Springer Verlag, 1996.
- Müller, J.P., M. Pischel, and M. Thiel (1995). Modelling reactive behavior in vertically layered agent architectures. In: (Wooldridge and Jennings, 1995a), pp. 261-276.
- Nwana, H.S., (1996), Software Agents: an Overview, *Knowledge Engineering Review*, vol. 11(3), pp. 205 - 244.
- Nwana, H.S., and Ndumu, D.T. (1998), A Brief Introduction to Software Agent Technology. In: (Jennings and Wooldridge, 1998b), pp. 29 – 47.
- Odell, J., Van Dyke Parunak, H., Bauer, B. (2000). Representing Agent Interaction Protocols in UML. *Proceedings AOSE 2000*. Springer Verlag, LNCS, vol. 1957, pp. 121-140
- Odell, J., Van Dyke Parunak, H., Fleischer, M. (2003a). Modeling Agents and their Environment: The Communication Environment. *Journal of Object Technology* 2(1): 39-52 (2003)
- Odell, J., Van Dyke Parunak, H., Fleischer, M. (2003b). Modeling Agents and their Environment: The Communication Environment. *Journal of Object Technology* 2(3): 39-52 (2003)
- Parunak, H. V. D. (1997). 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, vol. 75, pp. 69-101.
- Rao, A.S. and Georgeff, M.P. (1991). Modeling rational agents within a BDI architecture. In: R. Fikes and E. Sandewall (eds.), *Proceedings of the Second Conference on Knowledge Representation and Reasoning*, Morgan Kaufman, pp. 473-484.
- Rizzi, M., Baltés, M., Theobald, U. & Reuss, M. (1997). *In vivo* analysis of metabolic dynamics in *Saccharomyces cerevisiae*: II. Mathematical model. *Biotechnol. Bioeng.* **55**, 592–608.
- Rohwer, J.M., Meadow, N.D., Roseman, S., Westerhoff, H.V., Postma, P.W. (2000). Understanding glucose transport by the bacterial phosphoenolpyruvate:glycose phosphotransferase system on the basis of kinetic measurements in vitro. *J. Biol. Chem.*, 275(45):34909-21.
- Rotterdam, B.J. van, Crielaard, W., van Stokkum, I.H., Hellingwerf, K.J., Westerhoff H.V. (2002). Simplicity in complexity: the photosynthetic reaction center performs as a simple 0.2 V battery. *FEBS Lett.*, 510(1-2):105-7.
- Sharpankykh, A., and Treur, J. (2006). Verifying Interlevel Relations within Multi-Agent Systems. In: *Proceedings of the 17th European Conference on Artificial Intelligence*, Riva del Garda. IOS Press, 2006.
- Stuart, A.M., and Humphries, A.R. (1996). *Dynamical Systems and Numerical Analysis*. Cambridge University Press.
- Teusink, B., Passarge, J., Reijenga, C.A., Esgalhadó, E., van der Weijden, C.C., Schepper, M., Walsh, M.C., Bakker, B.M., van Dam, K., Westerhoff, H.V., Snoep, J.L. (2000). Can yeast

glycolysis be understood in terms of in vitro kinetics of the constituent enzymes? Testing biochemistry. *Eur J Biochem.*, 267(17):5313-29.

Vauclair, J. (1996). *Animal Cognition*. Harvard University Press, Cambridge, Massachusetts.

Westerhoff, H.V. (2001) The silicon cell, not dead but live! *Metab. Eng.* 2001; 3(3):207-10.

Wooldridge, M., Jennings, N.R. (eds.) (1995a) *Intelligent Agents*, Lecture Notes in Artificial Intelligence, Vol. 890, Springer Verlag, Berlin.

Wooldridge, M.J., Jennings, N.R. (1995b) Agent theories, architectures, and languages: a survey. In: (Wooldridge and Jennings, 1995a), pp. 1-39.

Wooldridge, M.J., and Jennings, N.R. (1995c). Intelligent Agents: Theory and practice. In: *Knowledge Engineering Review*, 10(2), pp. 115-152.