

Simulation and Analysis of Shared Extended Mind

Tibor Bosse¹, Catholijn M. Jonker¹, Martijn C. Schut¹, and Jan Treur^{1,2}

¹Vrije Universiteit Amsterdam, Department of Artificial Intelligence

De Boelelaan 1081a, NL-1081 HV Amsterdam, The Netherlands

Email: {tbosse, jonker, schut, treur}@cs.vu.nl

URL: <http://www.cs.vu.nl/~{tbosse, jonker, schut, treur}>

²Utrecht University, Department of Philosophy

Heidelberglaan 8, 3584 CS Utrecht

Abstract Some types of animals exploit patterns created in the environment as external mental states, thus obtaining an extension of their mind. In the case of social animals the creation and exploitation of such patterns can be shared, which supports a form of shared extended mind or collective intelligence. This paper explores this shared extended mind principle for social animals in more detail. The focus is on formal analysis and formalisation of the dynamic properties of the processes involved, both at the local level (the basic mechanisms) and the global level (the emerging properties of the whole), and their relationships. A case study in social ant behaviour in which shared extended mind plays an important role is used as illustration. For this case simulations are described based on specifications of local properties, and global properties are specified and verified.

1 Introduction

In [3], [4], [5], [6], [13] it is described how behaviour is often not only supported by an internal mind in the sense of internal mental structures and cognitive processes, but also by processes based on patterns created in the external environment that serve as external mental structures. Examples of this pattern of behaviour are the use of ‘to do lists’ and ‘lists of desiderata’. Having written these down externally (e.g., on paper, in your diary, in your organiser or computer) makes it unnecessary to have an internal memory about all the items. Thus internal mental processing can be kept less complex. The only thing to remember is where these lists are available. Other examples of the use of extended mind are doing mathematics or arithmetic, where external (symbolic, graphical, material) representations are used; e.g., [2]. In [3], [4], [5], [6], [13], [14] forms of cognitive processes like these are discussed in more detail.

Clark and Chalmers [5] point at the similarity between cognitive processes in the head and some processes involving the external world. This similarity can be used as an indication that these processes can be considered extended cognitive processes or extended mind: ‘If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain’t (all) in the head!’ [5], Section 2. (...) ‘Of course, one could always try to explain my action in terms of internal processes

and a long series of "inputs" and "actions", but this explanation would be needlessly complex. If an isomorphic process were going on in the head, we would feel no urge to characterize it in this cumbersome way.' [5], Section 3. We will call this criterion the 'isomorphism' criterion. As the patterns in the external world have to be created and sensed, interaction with the external world will be more intensive, compared to the case where internal mental states are created and exploited.

Especially in the case of social animals external mental states created by one individual can be exploited by another individual, or, more general, the creation and maintenance, as well as the exploitation of external mental states can be activities in which a number of individuals participate. For example, presenting slides on a paper with multiple authors to an audience. Further examples can be found everywhere, varying from roads, and traffic signs to books or other media, and to many other kinds of cultural achievements. In this multi-agent case the extended mind principle serves as a way to build a form of social or collective intelligence, that goes beyond (and may even not require) social intelligence based on direct one-to-one communication. In such cases the external mental states cross, and in a sense break up, the borders between (the minds of) the individuals and become shared mental states.

An interesting and currently often studied example of collective intelligence is the intelligence shown by ant colonies [1], [7]. Indeed, in this case the external world is exploited as an extended mind by using pheromones. While they walk, ants drop pheromones on the ground. The same or other ants sense these pheromones and follow the route in the direction of the strongest concentration. Pheromones evaporate over time; therefore such routes can vary over time.

In this paper this shared mind principle is analysed in more detail, and a formalisation is provided of its dynamics. These are illustrated by a case study of social behaviour based on shared extended mind (a simple ant colony). The analysis of this case study comprises multi-agent simulation based on identified local dynamic properties, identification of dynamic properties for the overall process, and verification of these dynamic properties.

2 State Properties and Dynamic Properties

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 mental (state) properties that do or do not hold at a certain point in time. For example, the internal state property 'the agent A has pain', or the external world state property 'the environmental temperature is 7° C', may be expressed in terms of different ontologies. To formalise state property descriptions, an ontology is specified as a finite set of sorts, constants within these sorts, and relations and functions over these sorts. The example properties mentioned above then can be defined by nullary predicates (or proposition symbols) such as `pain`, or by using n-ary predicates (with $n \geq 1$) like `has_temperature(environment, 7)`. For a given ontology Ont , the propositional language signature consisting of all *state ground atoms* (or *atomic state properties*) based on Ont is denoted by $APROP(Ont)$. 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. A *state* S is an indication of which atomic state properties are true and which are false, i.e., a mapping $S: \text{APROP}(\text{Ont}) \rightarrow \{\text{true}, \text{false}\}$.

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. A simple example is the following informally stated dynamic property for belief creation based on observation:

‘if the agent observes at t_1 that it is raining, then the agent will believe that it is raining’.

To express such dynamic properties, and other, more sophisticated ones, the Temporal Trace Language TTL is used; cf. [9]. To express dynamic properties in a precise manner a language is used in which explicit references can be made to time points and traces. Here *trace or trajectory* over an ontology Ont is a time-indexed sequence of states over Ont . The sorted predicate logic temporal trace language TTL is built on atoms referring to, e.g., traces, time and state properties. For example, ‘in trace γ at time t property p holds’ is formalised by $\text{state}(\gamma, t) \models p$. Here \models is a predicate symbol in the language, usually used in infix notation, which is comparable to the Holds -predicate in situation calculus. Dynamic properties are expressed by temporal statements built using the usual logical connectives and quantification (for example, over traces, time and state properties). For example, consider the following dynamic property:

‘in any trace γ , if at any point in time t_1 the agent A observes that it is raining,

then there exists a time point t_2 after t_1 such that at t_2 in the trace the agent A believes that it is raining’.

In formalised TTL form it looks as follows:

$$\forall t_1 \ [\text{state}(\gamma, t_1) \models \text{observes}(A, \text{itsraining}) \Rightarrow \\ \exists t_2 \geq t_1 \ \text{state}(\gamma, t_2) \models \text{belief}(A, \text{itsraining}) \]$$

Language abstractions by introducing new (definable) predicates for complex expressions are possible and supported.

A simpler temporal language has been used to specify simulation models. This language (the *leads to* language) enables to model direct temporal dependencies between two state properties in successive states. This executable format is defined as follows. Let α and β be state properties of the form ‘conjunction of atoms or negations of atoms’, and e, f, g, h non-negative real numbers. In the *leads to* language $\alpha \xrightarrow{e, f, g, h} \beta$, means:

If state property α holds for a certain time interval with duration g ,

then after some delay (between e and f) state property β will hold for a certain time interval of length h .

For a precise definition of the *leads to* format in terms of the language TTL , see [10]. A specification of dynamic properties in *leads to* format has as advantages that it is executable and that it can often easily be depicted graphically.

3 Explanation and the Isomorphism Principle

In Figure 1 (if the vertical dotted arrows are ignored for a moment), the upper part depicts a simple case of an agent with behaviour based on an internal mental state property m_1 and the lower part depicts another agent with the same behaviour based on an external mental state property m_2 . For example, the internal (m_1) or external (m_2) state property can be considered an agent’s belief. A more elaborate model is

described in Section 5. Note that the internal processing of the second agent is chosen as simple as possible: stimulus response. Hence, this agent is assumed not to have any internal states. The interaction between this agent and the external world is a bit more complex: one extra action is needed to create the external mental state m2, and one additional observation is needed to observe it.

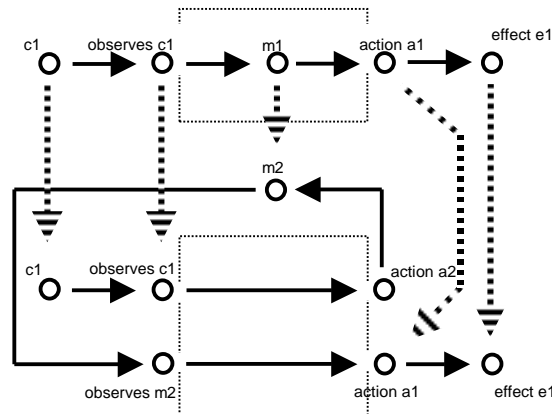


Figure 1 Internal and external mental states and their isomorphism relationship

To make the similarity between the two different cognitive processes more precise, the following mapping from the upper graph into the lower graph can be made:

c1	→	c1
observes c1	→	observes c1
m1	→	m2
action a1	→	action a1
effect e1	→	effect e1

This mapping, indicated by the vertical dotted arrows in Figure 1, preserves the temporal (leads to) relationships (the solid arrows) and provides an (isomorphic in the mathematical sense) embedding of a cognitive process based on internal mind into a cognitive process based on extended mind. Remember the quotes from [5], cited in the Introduction. Clark and Chalmers [5] use the isomorphism to a process ‘in the head’ as one of the criteria to consider external and interaction processes as cognitive, or mind processes. This ‘isomorphism’ criterion is formalised in Figure 1 for a simple example of such an isomorphism. Note that the process from m1 to action a1, modelled as one step in the internal case, is mapped onto a process from m2 via observes m2 to action a1, in the external case modelled as a two-step process. So the isomorphism is an embedding in one direction, not a bidirectional isomorphism, simply because the observation state for m2 (and the same for the action a2) has no counterpart in the internal case.

Behaviour often is explained by considering the basic underlying causal relations or mechanisms. Such basic mechanisms can be formally modelled by ‘leads to’ relations. The isomorphism principle and its formalisation as depicted in Figure 1 allows to replace an explanation of behaviour in terms of basic mechanisms involving frequent interactions (observations and actions) with the external world, by an explanation that leaves out these interactions and bases itself directly on the mental states.

This explanation is simpler, more abstract and perhaps more elegant, than the more complicated ‘cumbersome’ explanation based on the interactions. This is made possible by introducing a new ontology for the external world states involved. Considering part of the external world as extended mind allows one to give another interpretation to external physical processes and states. Physical state properties such as ‘pheromone is present at d’ are renamed as, for example, ‘it is believed that d is the direction home’. Why would one introduce extra language to refer to the same fact in the world? Given the literature on reduction, where often it is claimed that mental state properties can be and actually should be replaced by their physical realisers, at first sight such an opposite move may seem a bit surprising. For example, Kim [12] (pp. 214-216) claims that ontological simplification is one of the reasons to reduce mental state properties to physical state properties. In the extended mind case at hand the converse takes place; a question is what is the advantage of this ontological complication. A number of arguments in support of this can be given. By Clark and Chalmers [5], it is claimed that this allows application of other types of explanation and other methods of scientific investigation:

‘... we allow a more natural explanation of all sorts of actions. (...) in seeing cognition as extended one is not merely making a terminological decision; it makes a significant difference to the methodology of scientific investigation. In effect, explanatory methods that might once have been thought appropriate only for the analysis of "inner" processes are now being adapted for the study of the outer, and there is promise that our understanding of cognition will become richer for it.’ [5], Section 3.

In [11] it is explained in some detail, and illustrated by examples, why in various cases in other areas (such as Computer Science) such an antireductionist strategy often pays off; some of the discussed advantages in terms of insight, transparency and generality are: additional higher-level ontologies can improve understanding as they may allow simplification of the picture by abstracting from lower-level details; more insight is gained from a conceptually higher-level perspective; analysis of more complex processes is possible; the same concepts have a wider scope of application, thus obtaining unification. For more details and support for this antireductionist argument, see [11].

4 A Simulation Model of Shared Extended Mind

Dynamic properties can be specified at different aggregation levels, varying from (local) dynamic properties for the basic mechanisms and (global) properties of a process as a whole. This section introduces the local dynamic properties for the basic mechanisms; they are used to specify a simulation model. The world in which the ants live is described by a labeled graph as depicted in Figure 2.




Figure 2 An ants world

Locations are indicated by A, B, ..., and edges by l1, l2, ... The ants move from location to location via edges; while passing an edge, pheromones are dropped. The objective of the ants is to find food and bring this back to their nest. In this example there is only one nest (at location A) and one food source (at location F).

To model the example a number of local dynamic properties are used, some of which are presented below. Within these properties the following predicates are used:

connected_to_via(x0,x1,l)	location x0 is connected to x1 by edge l
neighbours(x, i)	location x has i neighbours
nest_location(x)	location x is the nest location
food_location(x)	location x is the food location
pheromones_at(l, i)	the pheromone level at edge l is i
arrives_at_location_from(a, x, l)	agent a arrives at location x from edge l
is_dropping_pheromones_at_from(a, l, x)	agent a drops pheromones at edge e, coming from location x
is_carrying_food(a)	agent a is carrying food
is_dropping_food(a)	agent a is dropping food

LP6 (Path Selection and Dropping Pheromones)

This property expresses that at a location where two different continuations (edges) with pheromones are possible, the edge with the highest pheromone level is taken and pheromones are dropped at this edge. Formalisation:

arrives_at_location_from(a, x0,l0) and neighbours(x0,3) and connected_to_via(x0,x1,l1) and pheromones_at(l1,i1) and connected_to_via(x0,x2,l2) and pheromones_at(l2,i2) and l0≠l1 and l0≠l2 and l1≠l2 and i1>i2 → is_dropping_pheromones_at_from(a,l1,x0)

For situations where no pheromone is present (yet) on the edges, the choice where to go is based on a concept of attractiveness of an edge (for example based on how difficult to go), represented by a predicate *attractive_direction_at(a, x, l)*. Some other local properties express that dropped pheromones are added to the present pheromone level; moreover, the amount of evaporated pheromone is subtracted. After passing an edge an ant arrives at a next location:

LP7 (Arrival at Next Location)

After passing an edge l, coming from location x1, and while dropping pheromones, an ant arrives at the next location x2. Formalisation:

is_dropping_pheromones_at_from(a,l,x1) and connected_to_via(x1,x2,l) → arrives_at_location_from(a,x2,l)

LP9 (Picking Up Food)

Arriving at a food location F food is picked up. Formalisation:

arrives_at_location_from(a,x,l) and food_location(x) → is_carrying_food(a)

LP10 (Dropping Food)

Arriving at the nest, while carrying food, an ant is dropping the food. Formalisation:

arrives_at_location_from(a,x,l) and nest_location(x) and is_carrying_food(a)
→ is_dropping_food(a)

A special software environment has been created to enable the simulation of executable models. Based on an input consisting of dynamic properties in *leads to* format, the software environment generates simulation traces. An example of such a trace can be seen in Figure 3. Time is on the horizontal axis, the state properties are on the vertical axis. A dark box on top of the line indicates that the property is true during that time period, and a lighter box below the line indicates that the property is false. This trace is based on all local properties identified. Because of space limitations, in the example depicted in Figure 3, only three ants are involved. However,

similar experiments have been performed with a population of 50 ants. Since the abstract way of modelling used for the simulation is not computationally expensive, also these simulations took no more than 30 seconds.

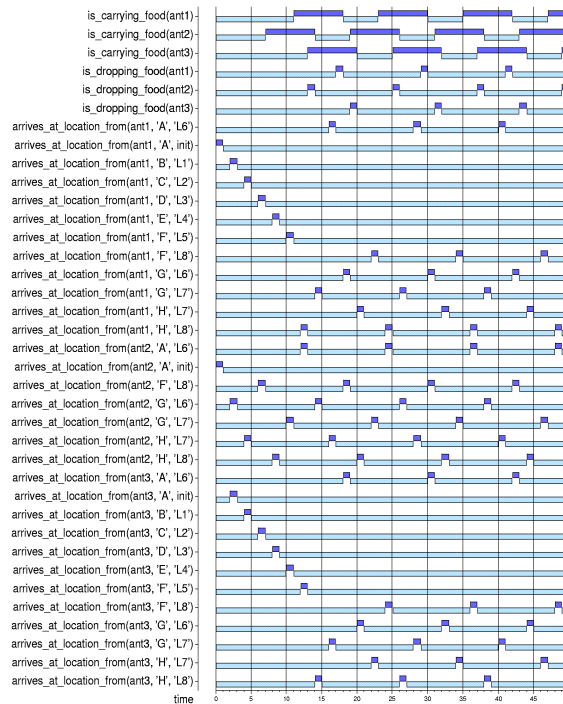


Figure 3 Simulation trace of the dynamics of the ants behaviour

As can be seen in Figure 3, there are two ants (ant1 and ant2) that start their search for food immediately, whereas ant3 comes into play a bit later, at time point 2. When ant1 and ant2 start their search, none of the locations contain any pheromones yet, so basically they have a free choice where to go. In the current example, ant1 selects a rather long route to the food source (via locations A-B-C-D-E-F), whilst ant2 chooses a shorter route (A-G-H-F). Note that, in the current model, a fixed route preference (via the attractiveness predicate) has been assigned to each ant for the case there are no pheromones yet. After that, at time point 3, ant3 starts its search for food. At that moment, there are trails of pheromones leading to both locations B and G, but these trails have exactly the same strength. Thus, ant3 also has a free choice among location B and G, and chooses in this case to go to B. From that moment on, ant3 will always be some steps behind ant1. Consequently, it will keep on following the trail of ant1. Meanwhile, at time point 6, ant2 has arrived at the food source (location F). Since it is the first to discover this location, the only present trail leading back to the nest, is its own trail. Thus ant2 will return home via its own trail. Next, when ant1 discovers the food source (at time points 10), it will notice that there is a trail leading back that is stronger than its own trail (since ant2 has already walked there twice: back and forth, not too long ago). As a result, it will also follow this trail. The same holds even

stronger for ant3. To conclude, eventually the shortest of both routes is shown to remain, whilst the other route evaporates. Other simulations, in particular for small ant populations, show that it is important that the decay parameter of the pheromones is not too high. Otherwise, the trail leading to the nest has evaporated before the first ant has returned, and all ants get lost!

5 Global Properties and Verification

In the previous section dynamic properties at the lowest aggregation level were addressed, the local dynamic properties, and simulation based on these properties was discussed. The current section addresses dynamic properties of a global nature, and their verification. Within these properties, γ is a variable that stands for an arbitrary trace. First a language abstraction is given:

$$\text{food_delivered_by}(\gamma, t, a) \equiv \exists x, l \ [\text{state}(\gamma, t) \models \text{arrives_at_location_from}(a, x, l) \ \& \ \text{state}(\gamma, t) \models \text{nest_location}(x) \ \& \ \text{state}(\gamma, t) \models \text{is_dropping_food}(a)]$$

GP1 Food Delivery Successfulness

There is at least one ant that brings food back to the nest.

$$\exists t \exists a: \text{food_delivered_by}(\gamma, t, a).$$

GP2 Multiple Delivery

Food is delivered by more than one ant

$$\exists t1, t2 \exists a1, a2 \ [\ a1 \neq a2 \ \& \ \text{food_delivered_by}(\gamma, t1, a1) \ \& \ \text{food_delivered_by}(\gamma, t2, a2) \]$$

Other language abstractions are:

$$\text{attractive_route_to}(\gamma, a, z) \equiv \exists x \exists l \ \forall t \ [\ \text{state}(\gamma, t) \models \text{attractive_direction_at}(a, x, l) \ \& \ \text{state}(\gamma, t) \models \text{connected_to_via}(x, z, l) \]$$

I.e., the attractive route of ant a passes through location z . For an explanation of the predicate $\text{attractive_direction_at}(a, x, l)$, see the previous section (below LP6).

$$\text{reaches_end_attractive_route}(\gamma, t, a) \equiv \exists x, l \ [\ \text{state}(\gamma, t) \models \text{arrives_at_location_from}(a, x, l) \ \& \ \text{attractive_route_to}(\gamma, a, x) \ \& \ \forall l' \ \text{state}(\gamma, t) \not\models \text{attractive_direction_at}(a, x, l') \]$$

GP3 Reaching End of Attractive Route

Ants reach the end of their attractive route.

$$\forall a \exists t \ \text{reaches_end_attractive_route}(\gamma, t, a)$$

GP4 Returning To Nest

Ants get back to the nest from the end of their attractive routes.

$$\forall a \forall t1 \exists l, t2 > t1 \exists x \ [\ \text{reaches_end_attractive_route}(\gamma, t1, a) \ \Rightarrow \ \text{state}(\gamma, t2) \models \text{arrives_at_location_from}(a, x, l) \ \& \ \text{state}(\gamma, t2) \models \text{nest_location}(x) \]$$

GP5 From Food To Nest

Ants get back to the nest from locations of food.

$$\forall a, l \ \forall t1 \exists t2 > t1 \exists z, x, l' \ [\ \text{state}(\gamma, t1) \models \text{arrives_at_location_from}(a, z, l) \ \& \ \text{state}(\gamma, t1) \models \text{food_location}(z) \] \ \Rightarrow \ \text{state}(\gamma, t2) \models \text{arrives_at_location_from}(a, x, l') \ \& \ \text{state}(\gamma, t2) \models \text{nest_location}(x)$$

These and a number of other global properties have been formalised and using a checking software environment have been (automatically) *verified in simulation traces*. This is a first manner for verification. A second way of verification is to es-

establish *logical relationships* between properties (by mathematical proof). This also has been performed in a number of cases. For example, under a number of assumptions the following relationships hold:

$GP4 \Rightarrow GP5$

$GP3 \ \& \ GP4 \Rightarrow GP2$

The assumptions include:

- attractive routes are not branching and are not crossing each other or themselves.
- at least two ants exist for which the attractive routes end at a food location and are short enough compared to the evaporation rate of pheromones to return.
- GP5 is only valid in the infinite future, since food sources are not depleted. In practice, the simulations stop, invalidating GP5 for the ants that are still on their way to the nest.

Furthermore, an additional premise Temporal Completion, see [8], is needed. For example, any of the following trivial (non-intended) world situations would disturb the ants: an ant comes to a location that contains a pheromone that is there without any reason (no ant dropped it), or on its way back an ant comes to a location without a pheromone (the pheromone immediately disappeared). It is clear that the above properties can only be proven under the assumption that nothing unexpected will happen. To put it differently, proofs can be given under the assumption that the set of local properties determines the whole range of events. This assumption has been added as a premise to establish the logical relationships between the properties.

6 Discussion

By Clark and Chalmers [5], Section 5, four criteria are given for extended mind: (1) the external information is a constant in the agent's life - when the information is relevant, he will rarely take action without consulting it; (2) the external information is directly available; (3) the agent endorses retrieved external information; (4) the external information has been endorsed at some point in the past, and is there as a consequence of this endorsement. In how far do these criteria apply to the ants case? For the first one, indeed an ant always senses the pheromone before choosing a direction. Secondly, at each location the pheromone is immediately accessible for sensing. Thirdly, the decision for the direction indeed is always based on the pheromone. Finally, the external information is endorsed in the past: the pheromone was dropped at the direction from which one or more ants came.

The extended mind perspective introduces an additional, cognitive ontology to describe properties of the physical world, which essentially is an antireductionist step, providing a more abstract and better manageable, higher level conceptualisation. In [11] a number of arguments can be found of why such antireductionist steps can be useful in explanation and theory development; also see Section 3 above. Indeed, following the extended mind perspective a high-level conceptualisation was obtained. This high-level conceptualisation could be formalised and analysed in a logical manner. The formalisation enables simulation and automated checking of dynamic properties of traces or sets of traces, and logically relating dynamic properties of different

aggregation levels to each other. All this would have been more difficult in case of an algorithmic or physically-oriented modelling perspective, involving, for example, differential equations and gradients of concentrations.

Work is currently in progress to replace the behaviour prescribed by attractiveness of a route by random route selection. In addition, experiments with food sources at different distances from the nest will be undertaken to determine the relation between evaporation rate and ants finding their way home. Therefore, these food sources will be made depletive. Also, the effect will be studied of using different types of pheromones. Moreover, an advanced visualisation environment is currently developed to make the simulation traces more readable. Finally, work is in progress to elaborate the isomorphism principle mentioned in Section 3 in more detail.

References

1. Bonabeau, J. Dorigo, M. and Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
2. Bosse, T., Jonker, C.M., and Treur, J. (2002). Simulation and analysis of controlled multi-representational reasoning processes. *Proc. of the Fifth International Conference on Cognitive Modelling, ICCM'03*. Universitäts-Verlag Bamberg, 2003, pp. 27-32.
3. Clark, A. (1997). *Being There: Putting Brain, Body and World Together Again*. MIT Press, 1997.
4. Clark, A. (2001). Reasons, Robots and the Extended Mind. In: *Mind & Language*, vol. 16, 2001, pp. 121-145.
5. Clark, A., and Chalmers, D. (1998). The Extended Mind. In: *Analysis*, vol. 58, 1998, pp. 7-19.
6. Deneubourg, J.L., Aron S, Goss S., Pasteels J. M. and Duerinck G. (1986). Random Behavior, Amplification Processes and Number of Participants: How They Contribute to the Foraging Properties of Ants. In: *Evolution, Games and Learning: Models for Adaptation in Machines and Nature*, North Holland, Amsterdam, 1986, pp. 176-186.
7. Dennett, D.C. (1996). *Kinds of Mind: Towards an Understanding of Consciousness*, New York: Basic Books.
8. Engelfriet, J. Jonker, C.M., and Treur, J. (2002). Compositional verification of Multi-Agent Systems in Temporal Multi-Epistemic Logic. *Journal of Logic, Language and Information*, vol. 11, 2002, pp. 195-225.
9. Jonker, C.M., and Treur, J., 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.
10. Jonker, C.M., Treur, J., and Wijngaards, W.C.A., A Temporal Modelling Environment for Internally Grounded Beliefs, Desires and Intentions. *Cognitive Systems Research Journal*, vol. 4(3), 2003, pp. 191-210.
11. Jonker, C.M., Treur, J., and Wijngaards, W.C.A., Reductionist and Antireductionist Perspectives on Dynamics. *Philosophical Psychology Journal*, vol. 15, 2002, pp. 381-409.
12. Kim, J. (1996). *Philosophy of Mind*. Westview Press
13. Kirsh, D. & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, vol. 18, 1994, pp. 513-49.
14. Menary, R. (ed.) (2004). *The Extended Mind*, Papers presented at the Conference *The Extended Mind - The Very Idea: Philosophical Perspectives on Situated and Embodied Cognition*, University of Hertfordshire, 2001. John Benjamins, 2004, to appear.