

Formal Interpretation of a Multi-Agent Society as a Single Agent

Tibor Bosse and Jan Treur

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

Abstract. In this paper the question is addressed to what extent the collective processes in a multi-agent society can be interpreted as single agent processes. This question is answered by formal analysis and simulation. It is shown for an example process how it can be conceptualised, formalised and simulated in two different manners: from a single agent (or cognitive) and from a multi-agent (or social) perspective. Moreover, it is shown how an ontological mapping can be formally defined between the two formalisations, and how this mapping can be extended to a mapping of dynamic properties. Thus it is shown how collective behaviour can be interpreted in a formal manner as single agent behaviour.

1 Introduction

Many processes in the world can be conceptualised using an *agent* metaphor. The result of such a conceptualisation is either a *single agent* (or cognitive) description or a *multi-agent* (or social) description. Especially for processes that are distributed, it is natural to describe them as a group of interacting agents. If a group of agents acts in a coherent way, however, one is often tempted to intuitively and informally interpret the process in singular form as a collective, and, in fact, as one individual (super)agent. The question addressed in this paper is whether in certain cases such an informal interpretation of a multi-agent system, acting in a collective manner, as an individual can be supported by a formal analysis. The approach to address this question is by formally defining an interpretation mapping between a conceptualisation of a process as a multi-agent system and a conceptualisation of the same process as an individual.

The prerequisites to undertake such a formal analysis concern formalisations of the notion of agent, single agent behaviour and multi-agent behaviour, and the notion of interpretation mapping. More specifically, what is needed is a formal notion of what an agent is in the sense of

- distinctions between the agent's internal mental processes, the agent's body, and the agent's environment
- interactions and relationships between mental aspects and body aspects
- interactions and relationships between agent and environment, including interactions with other agents

Furthermore, formalisations of single agent behaviour and multi-agent behaviour are needed that cover

- the externally observable behaviour
- the underlying internal processes

Moreover, a formal notion of interpretation mapping of a single agent conceptualisation into a multi-agent conceptualisation is needed that

- maps ontological concepts describing a conceptualisation of a process from an individual perspective to ontological concepts describing a conceptualisation of the same a process from a multi-agent perspective
- covers mapping of individual mental state properties for the single agent conceptualisation to shared mental state properties for the multi-agent conceptualisation
- covers the mapping of dynamic aspects of single agent behaviour onto those of multi-agent behaviour

In this paper, for these three notions formalisations are provided and used to indeed achieve an approach of how a collective can be formally interpreted as an individual.

The formal interpretation approach is evaluated for the case of collective behaviour of an ant colony. The intelligence shown by ant colonies are an interesting and currently often studied example of collective intelligence (Bonabeau et al. 1999; Deneubourg et al. 1986; Drogoul et al. 1995). In this case, by using pheromones, the external world is exploited as a form of extended mind; cf. (Clark 1997; Clark and Chalmers 1998; Dennett 1996; Kirsch and Maglio 1994; Menary 2006). It is shown (in Section 7) how this case can be seen as a paradigmatic case, also covering cases in human society. For example, cases in which an organization or department wants to interact with ‘one face’ with the outside world, and to this end maintains a repository for common guidelines. The analysis of this case study comprises on the one hand a multi-agent model, simulation based on identified local dynamic properties, and identification of dynamic properties for the overall process. On the other hand the same is done for an alternative model based on a single agent with internal mental states, and the two models are related to each other via the interpretation mapping.

In Section 2, a formalisation of basic agent concepts is introduced. Section 3 explains, using a simple example, the idea of the basic formal ontology mapping between state properties in a single agent conceptualisation and state properties in a multi-agent conceptualisation. In Section 4 this notion of basic interpretation mapping of state properties is applied to two conceptualisations of the more complex ant colony example, the central case study in the paper. Section 5 discusses the dynamics for the two conceptualisations of the ant colony example in more detail, which leads to formal specification of executable local dynamic properties that have been used for simulation. In Section 6 the basic interpretation mapping for state properties is extended to dynamic properties, thus obtaining an interpretation mapping between the two conceptualisations of the dynamics of the example ant colony process. In Section 7 it is shown how the interpretation approach can be applied to other types of societies (e.g., human societies), where patterns occur that are similar to those in the ants case. Section 8 is a final discussion.



2 Basic Agent Concepts

The agent perspective entails a distinction between the following different types of ontologies:

- an ontology for *internal mental properties* of the agent A ($MentOnt(A)$),
- for properties of the agent’s (physical) *body* ($BodyOnt(A)$),
- for properties of the (sensory or communication) *input* ($InOnt(A)$)
- for properties of the (action or communication) *output* ($OutOnt(A)$) of the agent, and
- for properties of the *external world* ($ExtOnt(A)$).

For example, the property ‘the agent A feels pain’ may belong to $\text{MentOnt}(A)$, resp. $\text{BodyOnt}(A)$, whereas ‘it is raining’ and ‘the outside temperature is 7°C ’ may belong to $\text{ExtOnt}(A)$. The agent input ontology InOnt defines state properties for received perception or communication, as an in-between step from environment or body state properties to internal mental state properties, the agent output ontology OutOnt defines state properties that indicate initiations of actions or communications of the agent, as an in-between step from internal mental state properties to environment or body state properties. The combination of InOnt and OutOnt is the *agent interaction ontology*, defined by $\text{InteractionOnt} = \text{InOnt} \cup \text{OutOnt}$.

To formalise state property descriptions of the types introduced above, 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. The example properties mentioned above then can be defined by nullary predicates (or proposition symbols) such as itsraining , or by using n-ary predicates (with $n \geq 1$) like $\text{has_pain}(A)$ and $\text{has_temperature}(\text{environment}, 7)$.

For a given ontology Ont , the propositional language signature consisting of all *state ground atoms* based on Ont is denoted by $\text{APROP}(\text{Ont})$. The *state properties* based on a certain ontology Ont are formalised by the propositions that can be made, using (using conjunction, negation, disjunction, implication) from the ground atoms. 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. In other words, 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 externally observable dynamics of the agent, explicit reference is made to time. Dynamics will be described as evolution of states over 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 sorted predicate logic *Temporal Trace Language* (TTL) is used (Jonker et al. 2003). Here, a *trace* over an ontology Ont is a time-indexed sequence of states over Ont . 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, the dynamic property put forward above can be expressed in a more structured semiformal manner as:

‘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:

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

Based on TTL, a simpler temporal language has been defined to specify simulation models. This language (the *leads to* language) enables modelling 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 $\alpha \bullet \rightarrow_{e, f, g, h} \beta$, 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 (Jonker et al. 2003). 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 The Basic Interpretation Mapping

In this section it is discussed how a conceptualisation based on a single agent and individual (internal) mental state properties can formally be mapped onto a conceptualisation based on multiple agents and shared (for the sake of simplicity assumed external) mental state properties. Here this ontological mapping is only given in its basic form, for the state properties. In Section 6 the basic mapping is extended to temporal expressions describing behaviour.

First, consider Figure 1. This figure depicts a simple case of a single agent A with behaviour based on an *individual internal* mental state property m_1 . The solid arrows depict temporal *leads to* relationships. Mental state property m_1 (temporally) depends on observations of three world state properties c_1 , c_2 , c_3 . Moreover, action a_1 depends on m_1 .

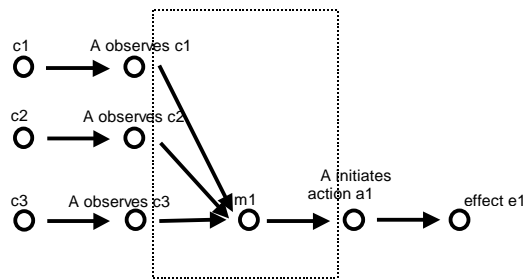


Figure 1. Single Agent behaviour based on an internal mental state

Now consider Figure 2. This figure depicts a group of agents A_1 , A_2 , A_3 , A_4 with behaviour based on a physical external world state property m_2 that serves as a *shared external* mental state property.

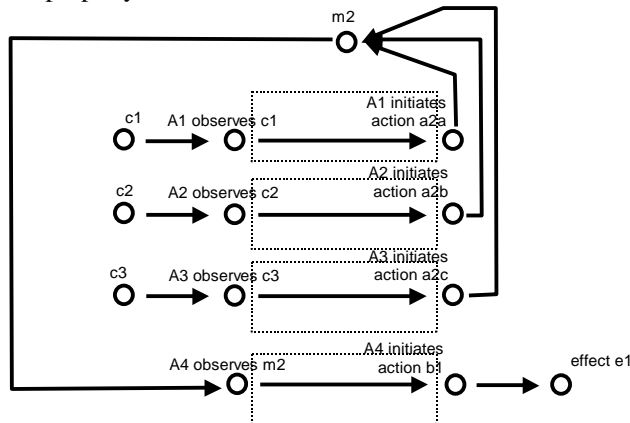


Figure 2. Multi-Agent behaviour based on a shared external mental state

To create this shared mental state property, actions a2a, a2b, a2c of the agents A1, A2, A3 are needed, and to show the behaviour, first an observation of m2 by agent A4 is needed. Note that here the internal processing is chosen as simple as possible: stimulus response. Hence, this agent is assumed not to have any internal states. This is in line with the ideas of Clark and Chalmers, who claim that the explanation of cognitive processes should be as simple as possible (Clark and Chalmers 1998). However, the interaction between agent and external world is a bit more complex: compared to a single agent perspective with internal mental state m1, extra actions of some of the agents needed to create the external mental state property m2, and additional observations are needed to observe it.

To make the similarity between the two different cognitive processes more precise, the following mapping from the nodes (state properties) in Figure 1 onto nodes in Figure 2 can be made (see Figure 3):

External world state properties

\varnothing : c1 → c1
 \varnothing : c2 → c2
 \varnothing : c3 → c3
 \varnothing : effect e1 → effect e1

Observation state properties

\varnothing : A observes c1 → A1 observes c1
 \varnothing : A observes c2 → A2 observes c2
 \varnothing : A observes c3 → A3 observes c3

Action initiation state properties

\varnothing : A initiates action a1 → A4 initiates action b1

Mental state property to external world state property

\varnothing : m1 → m2

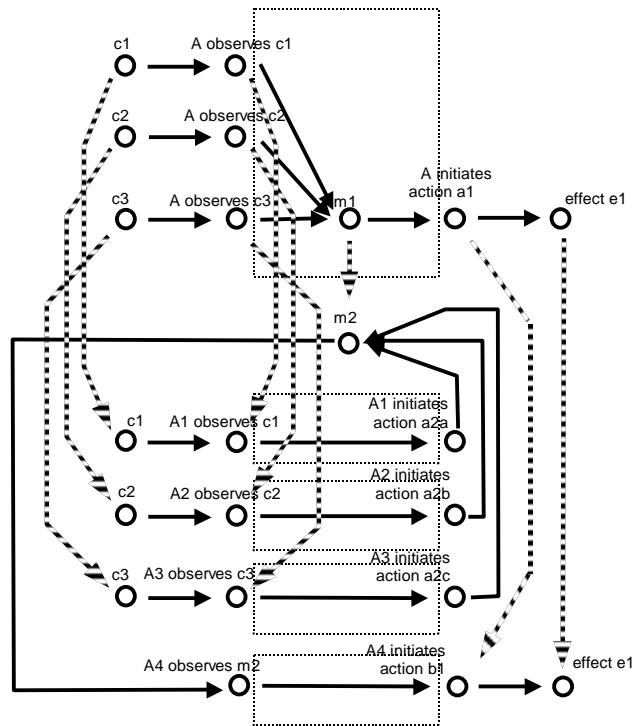


Figure 3. Mapping from individual mental state to shared extended mind

Note that in this case, for simplicity it is assumed that each observation of A is an observation of exactly one of the A_i , and the same for actions.

This mapping ϕ , indicated by the vertical dotted arrows in Figure 3, preserves the temporal dependencies in the form of *leads to* relationships (the solid arrows) and provides an isomorphic embedding (in the mathematical sense) of a cognitive process based on internal mind into a cognitive process based on extended mind.

In their paper about extended mind, Clark and Chalmers (1998) 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! (...)
(Clark and Chalmers 1998), Section 2.

One can explain my choice of words in Scrabble, for example, as the outcome of an extended cognitive process involving the rearrangement of tiles on my tray. 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. (...) In a very real sense, the re-arrangement of tiles on the tray is not part of action; it is part of *thought*.
(Clark and Chalmers 1998), Section 3.

Clark and Chalmers (1998) use the isomorphic relation to a process 'in the head' as one of the criteria to consider external and interaction processes as cognitive, or mind processes. As the shared mental state property m_2 is modelled as an external state property, this principle is formalised in Figure 3. Note that the process from m_1 to action a_1 , modelled as one step in the single agent, internal case, is mapped onto a process from m_2 via A_4 observes m_2 to A_4 initiates action b_1 , modelled as a two-step process in the multi-agent, external case. So the mapping is an isomorphic embedding in one direction, not a bidirectional isomorphism, simply because on the multi-agent side, the observation state for A_4 observing m_2 has no counterpart in the single agent, internal case (and the same for the agents A_1 , A_2 , A_3 initiating actions a_{2a} , a_{2b} , a_{2c}).

Notice that the mapping ϕ is a (formal) mapping between state properties. However, it was already put forward that temporal *leads to* relations are preserved under ϕ , so the mapping can be extended to a mapping of *leads to* properties onto *leads to* properties. From a more general perspective, it can be analysed how far the mapping ϕ can be extended to a (formal) mapping from dynamic properties to dynamic properties expressed in TTL. This will be addressed in detail in Section 6.



4 Two Conceptualisations and their Mapping

The general formalisation perspective put forward in previous sections has been evaluated for a case study: a process of collective ant behaviour. For this example process, two conceptualisations have been made, one from a multi-agent (or social) perspective, and one for a single agent (or cognitive) perspective. In Section 7 it is shown how this case can be seen as a paradigmatic case for a large class of cases, including cases in human society (for example, where an organization or department wants to behave with 'one face' to the external world by maintaining common guidelines).

The world in which the ants live is described by a labeled graph as depicted in Figure 4. Locations are indicated by A, B, \dots , and edges by E_1, E_2, \dots . To represent such a graph the predicate `connected_to_via(l_0, l_1, e_1)` is used. The ants move from location to location via edges; while passing an edge, pheromones are dropped. 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. The goal of the ants is to find food and bring this back to their nest. In this example there is only one nest (location A) and one food source (location F).

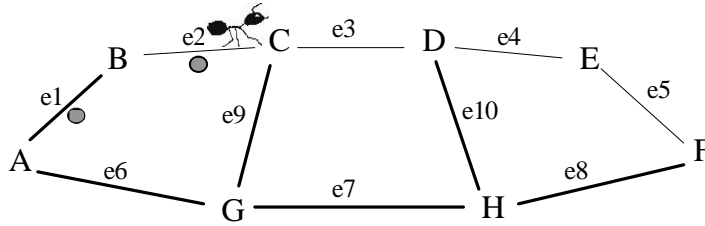


Figure 4. An ant world

4.1 Multi-Agent Conceptualisation

The example process conceptualised from a multi-agent perspective concerns multiple agents (the ants), each of which has input (to observe) and output (for moving and dropping pheromones) states, and a physical body which is at certain positions over time. However, following the claims in the previous section, they do not have any internal mental state properties (they are assumed to act purely by stimulus-response behaviour). Note that the reason for leaving out internal states is not that it is impossible, but simply that they are not needed for our purposes here. However, as will be discussed in Section 7, the interpretation approach is applicable to agents with internal states as well. An overview of the formalisation of the state properties of the multi-agent conceptualisation is shown in Table 1.

Table 1. Multi-Agent conceptualisation: state properties

	Multi-Agent Conceptualisation
	<i>body positions in world:</i>
pheromone level at edge e is i	pheromones_at(e, i)
ant a is at location l coming from e	is_at_location_from(a, l, e)
ant a is at edge e to l2 coming from location l1	is_at_edge_from_to(a, e, l1, l2)
ant a is carrying food	is_carrying_food(a)
	<i>world state properties:</i>
edge e connects location l1 and l2	connected_to_via(l1, l2, e)
location l has i neighbours	neighbours(l, i)
edge e is most attractive for ant a coming from location l	attractive_direction_at(a, l, e)
	<i>input state properties:</i>
ant a observes that it is at location l coming from edge e	observes(a, is_at_location_from(l, e))
ant a observes that it is at edge e to l2 coming from location l1	observes(a, is_at_edge_from_to(e, l1, l2))
ant a observes that edge e has pheromone level i	observes(a, pheromones_at(e, i))
	<i>output state properties:</i>
ant a initiates action to go to edge e to l2 coming from location l1	to_be_performed(a, go_to_edge_from_to(e, l1, l2))
ant a initiates action to go to location l coming from edge e	to_be_performed(a, go_to_location_from(l, e))
ant a initiates action to drop pheromones at edge e coming from location l	to_be_performed(a, drop_pheromones_at_edge_from(e, l))
ant a initiates action to pick up food	to_be_performed(a, pick_up_food)
ant a initiates action to drop food	to_be_performed(a, drop_food)

4.2 Single-Agent Conceptualisation

The conceptualisation of the example process from a single agent perspective (Superant S), however, takes into account one body, of which each ant is part (for convenience we call them the ‘paws’ of this body). Also the pheromone levels at the edges are part of the body.

Table 2. Single Agent conceptualisation: state properties

Single Agent Conceptualisation	
<i>mental state properties:</i>	
belief(S, relevance_level(e, i))	belief on the relevance level i of an edge e
<i>body position in world:</i>	
has_paw_at_location_from(S, p, l, e)	position of paw p at location l coming from edge e
has_paw_at_edge_from_to(S, p, e, l1, l2)	position of paw p at edge e to l2 coming from location l1
is_carrying_food_with_paw(S, p)	paw p is carrying food
<i>world state properties:</i>	
connected_to_via(l1, l2, e)	edge e connects location l1 and l2
neighbours(l, i)	location l has i neighbours
attractive_direction_at(p, l, e)	edge e is most attractive for paw p coming from location l
<i>input state properties:</i>	
observes(S, has_paw_at_location_from(p, l, e))	S observes that paw p is at location l coming from edge e
observes(S, has_paw_at_edge_from_to(p, e, l1, l2))	S observes that paw p is at edge e to l2 coming from location l1
<i>output state properties:</i>	
to_be_performed(S, move_paw_to_edge_from_to(p, e, l1, l2))	S initiates action to move paw p from location l1 to edge e to l2
to_be_performed(S, move_paw_to_location_from(p, l, e))	S initiates action to move paw p from edge e to location l
to_be_performed(S, pick_up_food_with_paw(p))	S initiates action to pick up food with paw p
to_be_performed(S, drop_food_with_paw(p))	S initiates action to drop food with paw p

The body position of this agent in the world is defined by the collection of positions of each of the paws. Mental state properties for this single agent occur in the form of beliefs that a certain edge has a certain relevance level (realised in the body by the pheromone levels). Input of the single agent is defined by the collection of inputs of the ants at each of the paws. Output is defined by initiation of movements of one or more of the paws. Notice that in this case dropping pheromones is not an action, but an internal body process to create or update the proper beliefs by creating or updating their realisation in the body. An overview of the formalisation of the state properties of the multi-agent conceptualisation is shown in Table 2. Note that there S stands for the Superant.

4.3 Mapping between Conceptualisations

The two conceptualisations described in Sections 4.1 and 4.2 are two conceptualisations of one and the same example process. A concept in any of the two conceptualisations in principle has a one-to-one correspondence to an aspect of this example process which can be considered the informal semantics of the concept (in our case the concept is formalised); see the double arrows in Figure 5.

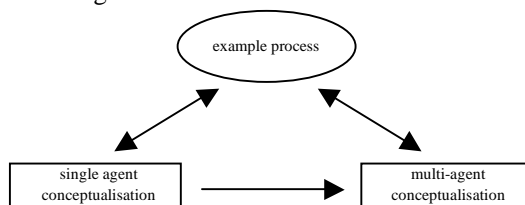


Figure 5. Two conceptualisations and their mapping

Given these one-to-one correspondences, a mapping from the single agent conceptualisation to the multi-agent conceptualisation can be made as follows:

- 1) Take any state property c belonging to the single agent conceptualisation
- 2) Identify to what aspect a of the example process this state property corresponds
- 3) Identify to which state property d in the multi-agent conceptualisation this aspect a corresponds
- 4) Map c to d .

If this approach works, then a mapping is obtained that is faithful with respect to the example process: the state property d to which c is mapped corresponds to the same aspect a of the process as c , and therefore will be true (for the informal semantics) if and only if c is. The approach can also fail. It can fail in 2) if state properties are used in the single agent conceptualisation that have no counterpart in the example process. It can fail in 3) if in the single agent conceptualisation, aspects of the process are covered that are left out of consideration in the other conceptualisation. Actually, such aspects exist the other way around: there are aspects of the process, such as observing the pheromones, that are covered by the multi-agent conceptualisation, but not by the single agent conceptualisation. Therefore such a mapping is not possible from right to left in Figure 5 (see also Figure 3 in Section 3, where the mapping is not bijective either). However, a mapping from left to right (single agent to multi-agent conceptualisation), is possible. It is shown in Table 3. Note that there S stands for the Superant, and paw p corresponds to ant a .

Table 3. Mapping between state properties

Single Agent Conceptualisation	Multi-Agent Conceptualisation
belief(S , relevance_level(e , i))	pheromones_at(e , i)
has_paw_at_location_from(S , p , l , e) has_paw_at_edge_from_to(S , p , e , $l1$, $l2$) is_carrying_food_with_paw(S , p)	is_at_location_from(a , l , e) is_at_edge_from_to(a , e , $l1$, $l2$) is_carrying_food(a)
connected_to_via($l1$, $l2$, e) neighbours(l , l) attractive_direction_at(p , l , e)	connected_to_via($l1$, $l2$, e) neighbours(l , i) attractive_direction_at(a , l , e)
observes(S , has_paw_at_location_from(p , l , e)) observes(S , has_paw_at_edge_from_to(p , e , $l1$, $l2$)) ---	observes(a , is_at_location_from(l , e)) observes(a , is_at_edge_from_to(e , $l1$, $l2$)) observes(a , pheromones_at(e , i))
to_be_performed(S , move_paw_to_edge_from_to(p , e , $l1$, $l2$)) to_be_performed(S , move_paw_to_location_from(p , l , e)) --- to_be_performed(S , pick_up_food_with_paw(p)) to_be_performed(S , drop_food_with_paw(p))	to_be_performed(a , go_to_edge_from_to(e , $l1$, $l2$)) to_be_performed(a , go_to_location_from(l , e)) to_be_performed(a , drop_pheromones_at_edge_from(e , l)) to_be_performed(a , pick_up_food) to_be_performed(a , drop_food)

5 Two Simulation Models

The two conceptualisations introduced above have been used to create two simulation models for collective ant behaviour: one from a multi-agent (social) perspective and one from a single agent (cognitive) perspective. The basic building blocks of the model were dynamic properties in *leads to* format, specifying the local mechanisms of the process. Examples of such local dynamic properties (for the *multi-agent case*) are the following:

LP5 (Selection of Edge)

“If an ant observes that it is at location l , and there are three edges connected to that location, then the ant goes to the edge with the highest amount of pheromones.” Formalisation:

observes(a, is_at_location_from(l, e0)) and neighbours(l, 3) and connected_to_via(l, l1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(l, l2, e2) and observes(a, pheromones_at(e2, i2)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2$ and $i1 > i2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, l1))

LP6 (Arrival at Edge)

“If an ant goes to edge e from location l to location l1, then later the ant will be at this edge e.”
to_be_performed(a, go_to_edge_from_to(e, l, l1)) \leftrightarrow is_at_edge_from_to(a, e, l, l1)

LP9 (Dropping of Pheromones)

“If an ant observes that it is at an edge e from a location l to a location l1, then it will drop pheromones at this edge e.”

observes(a, is_at_edge_from_to(e, l, l1)) \leftrightarrow to_be_performed(a, drop_pheromones_at_edge_from(e, l))

LP12 (Observation of Pheromones)

“If an ant is at a certain location l, then it will observe the number of pheromones present at all edges that are connected to location l.”

is_at_location_from(a, l, e0) and connected_to_via(l, l1, e1) and pheromones_at(e1, i) \leftrightarrow observes(a, pheromones_at(e1, i))

LP13 (Increment of Pheromones)

“If an ant drops pheromones at edge e, and no other ants drop pheromones at this edge, then the new number of pheromones at e becomes $i * \text{decay} + \text{incr}$.” Here, i is the old number of pheromones, decay is the decay factor, and incr is the amount of pheromones dropped.

to_be_performed(a1, drop_pheromones_at_edge_from(e, l1)) and $\forall l2$ not to_be_performed(a2, drop_pheromones_at_edge_from(e, l2)) and $\forall l3$ not to_be_performed(a3, drop_pheromones_at_edge_from(e, l3)) and $a1 \neq a2$ and $a1 \neq a3$ and $a2 \neq a3$ and pheromones_at(e, i) \leftrightarrow pheromones_at(e, $i * \text{decay} + \text{incr}$)

LP14 (Collecting of Food)

“If an ant observes that it is at location F (the food source), then it will pick up some food.” observes(a, is_at_location_from(F, e)) \leftrightarrow to_be_performed(a, pick_up_food)

To model the example from a single agent perspective, again a number of local dynamic properties are used. Most, but not all of these local properties have a 1:1 correspondence to those for the multi-agent case. For example, the properties for the *single agent case* that correspond to the properties above are as follows (see the next section for more information about this correspondence):

LP5' (Selection of Edge)

“If S observes that it has a paw p at location A, and there are three edges connected to that location, then S will move its paw to the edge of which it believes that it has the highest relevance level.”

observes(S, has_paw_at_location_from(p, l, e0)) and neighbours(l, 3) and connected_to_via(l, l1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(l, l2, e2) and belief(S, relevance_level(e2, i2)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2$ and $i1 > i2 \leftrightarrow$ to_be_performed(S, move_paw_to_edge_from_to(p, e1, l1))

LP6' (Paw Arrival at Edge)

“If S moves its paw p to an edge e from a location l to a location l1, then later this paw will be at this edge e.”

to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1)) \leftrightarrow has_paw_at_edge_from_to(S, p, e, l, l1)

LP11' (Increment of Belief)

“If S has exactly one paw at edge e, then the new number of pheromones at e becomes $i * \text{decay} + \text{incr}$.”

observes(S, has_paw_at_edge_from_to(p1, e, l, l1)) and $\forall l2$ not observes(S, has_paw_at_edge_from_to(p2, e, l, l2)) and $\forall l3$ not observes(S, has_paw_at_edge_from_to(p3, e, l, l3)) and $p1 \neq p2$ and $p1 \neq p3$ and $p2 \neq p3$ and belief(S, relevance_level(e, i)) \leftrightarrow belief(S, relevance_level(e, $i * \text{decay} + \text{incr}$))

LP12' (Collecting of Food)

“If S observes that it has a paw p at location F (the food source), then it will pick up some food with that paw.”

observes(S, has_paw_at_location_from(p, F, e)) \leftrightarrow to_be_performed(S, pick_up_food_with_paw(p))

The complete sets of local properties used to model the example are shown in Appendix A (multi-agent case) and Appendix B (single agent case).

In (Bosse et al. 2005), a special software environment is introduced for the simulation of executable models. Based on an input consisting of dynamic properties in *leads to* format, it can generate simulation traces. This environment has been used to generate a number of simulation traces for the ants case study. An example of (part of) such a trace can be seen in Figure 6. To facilitate understanding, in this simulation only three ants are involved. Moreover, only some of the relevant state properties are shown (in particular, those dealing with the movement of ant1, and with the food delivery of the ants). 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 was based on the multi-agent simulation model.

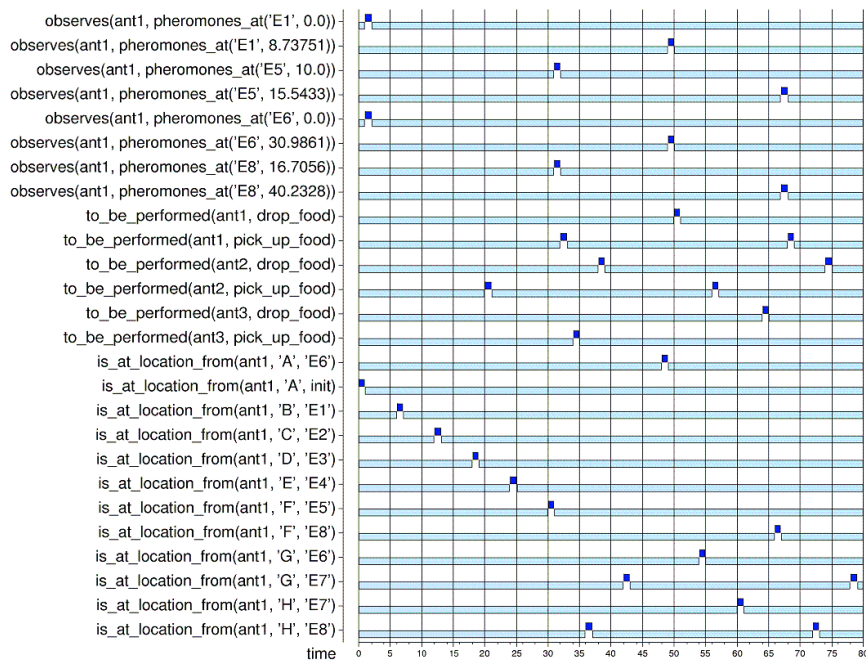


Figure 6. Multi-Agent Simulation Trace

Figure 7 depicts a similar trace as Figure 6, this time based on the single agent simulation model. Note that there are several differences between Figure 6 and 7. In the first place, all ants that are treated as separate agents in Figure 6, are considered as parts of Superant S in Figure 7. For example, *is_at_location_from(ant1, A, E6)* in the multi-agent case corresponds to *has_paw_at_location_from(S, paw1, A, E6)* in the single agent case. Another important difference is that in the single agent case, there is no explicit observation of pheromones. The reason for this is that the *belief(S, relevance_level(e, i))* states (which are the single agent equivalent for the *pheromones_at(e, i)* states in the multi-agent case) are internal states of S, which do not have to be observed.

Altogether, the software environment has been used to successfully generate a large number of simulation traces on the basis of both simulation models. As mentioned earlier, in the examples depicted only three ants are involved. However, similar experiments have

been performed with populations of 50 and 100 ants. Since the abstract way of modelling used for the simulation is not computationally expensive, also these simulations can be performed relatively quickly. To be precise, they took 35 seconds (for 50 ants and 80 time steps), 70 seconds (100 ants, 80 time steps), 100 seconds (50 ants, 200 time steps), and 200 seconds (100 ants, 200 time steps), respectively. A number of these simulation traces are stored in: <http://www.cs.vu.nl/~tbosse/isomorphism/>.

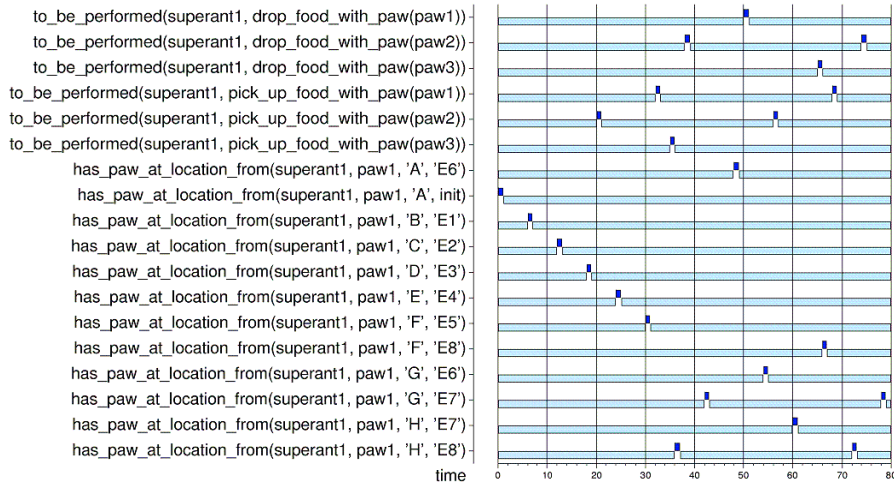


Figure 7. Single Agent Simulation Trace

6 The Extended Interpretation Mapping

In Section 3 it was shown how the basic interpretation mapping can be defined as a mapping between state properties. It was suggested that this mapping can be extended to a mapping between local dynamic properties in *leads to* format. Therefore, the following interpretation mapping can be defined:

$$\varphi(\alpha \rightarrow \beta) = \varphi(\alpha) \leftrightarrow \varphi(\beta)$$

Using this interpretation mapping, combined with the basic mapping of the state ontology elements described in Section 4, mappings between the dynamic properties of the case study can be found, e.g.:

```

φ(LP6')
= φ(to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1)) ↔ has_paw_at_edge_from_to(S, p, e, l, l1))
= φ(to_be_performed(S, move_paw_to_edge_from_to(p, e, l, l1))) ↔
  φ(has_paw_at_edge_from_to(S, p, e, l, l1))
= to_be_performed(a, go_to_edge_from_to(e, l, l1)) ↔ is_at_edge_from_to(a, e, l, l1)
= LP6

```

A mapping between all local dynamic properties (in *leads to* format) of the case study is given in Table 4. Notice that in some cases a certain dynamic property is mapped to a dynamic property that is not literally in the multi-agent model, but actually is a combination of two other local properties present in the model. This shows where the single agent conceptualisation is simpler than the multi-agent conceptualisation.

Table 4. Mapping between local dynamic properties

Single Agent Conceptualisation	Multi-Agent Conceptualisation
LP1'	LP1
LP2'	LP2
LP3'	LP3
LP4'	LP4
LP5'	LP5 & LP12
LP6'	LP6
LP7'	LP7
LP8'	LP8
LP9'	LP10
LP10'	LP11
LP11'	LP9 & LP13
LP12'	LP14
LP13'	LP15
LP14'	LP16
LP15'	LP17
LP16'	LP9 & LP18

The mapping shown in Table 4 is a syntactic mapping. However, also the traces generated on the basis of these properties can be mapped: each trace γ can be mapped onto a trace $\varphi(\gamma) = \gamma'$. For example, the trace depicted in Figure 7 can be mapped onto the trace depicted in Figure 6. This shows that the syntactic mapping between local properties preserves semantics.

In addition, it is possible to extend the mapping to the wider class of TTL expressions. Recall that TTL expressions are built on atoms of the form $state(\gamma, t) \models p$. By the basic mapping the state property p can be translated into $\varphi(p)$, which is assumed to be part of the ontology of one of the agents A_i in the multi-agent conceptualisation. Moreover, the trace name γ can be mapped onto a trace name $\varphi(\gamma) = \gamma'$. Then the extended interpretation mapping for $state(\gamma, t) \models p$ is defined by:

$$\varphi: state(\gamma, t) \models p = state(\gamma', t) \models \varphi(p)$$

After these atoms have been mapped, TTL expressions as a whole can be mapped in a straightforward compositional manner:

$$\begin{aligned} \varphi(A \& B) &= \varphi(A) \& \varphi(B) \\ \varphi(A \Rightarrow B) &= \varphi(A) \Rightarrow \varphi(B) \\ \varphi(\text{not } A) &= \text{not } \varphi(A) \\ \varphi(\forall v A(v)) &= \forall v' \varphi(A(v')) \\ \varphi(\exists v A(v)) &= \exists v' \varphi(A(v')) \end{aligned}$$

For example, take the following TTL expression, which is a global property for the single agent case of the ant example:

GP1' Food Discovery

‘Eventually, one of the paws of S will be at the food location.’

$$\exists t, p, l, e [state(\gamma, t) \models \text{has_paw_at_location_from}(S, p, l, e) \& state(\gamma, t) \models \text{food_location}(l)]$$

This expression is mapped as follows:

$$\varphi(\exists t, p, l, e [state(\gamma, t) \models \text{has_paw_at_location_from}(S, p, l, e) \& state(\gamma, t) \models \text{food_location}(l)])$$

$$\begin{aligned}
&= \exists t, p', l', e' \left(\varphi([\text{state}(\gamma, t) \models \text{has_paw_at_location_from}(S, p', l', e') \ \& \ \text{state}(\gamma, t) \models \text{food_location}(l')]) \right) \\
&= \exists t, p', l', e' \left(\varphi(\text{state}(\gamma, t) \models \text{has_paw_at_location_from}(S, p', l', e')) \ \& \ \varphi(\text{state}(\gamma, t) \models \text{food_location}(l')) \right) \\
&= \exists t, p', l', e' \left[\text{state}(\gamma, t) \models \varphi(\text{has_paw_at_location_from}(S, p', l', e')) \ \& \ \text{state}(\gamma, t) \models \varphi(\text{food_location}(l')) \right] \\
&= \exists t, p', l', e' \left[\text{state}(\gamma, t) \models \text{is_at_location_from}(p', l', e') \ \& \ \text{state}(\gamma, t) \models \text{food_location}(l') \right]
\end{aligned}$$

Thus, eventually global property GPI' is mapped to the following global property (GP1):

GP1 Food Discovery

"Eventually, one of the ants will be at the food location."

$$\exists t, a, l, e \left[\text{state}(\gamma, t) \models \text{is_at_location_from}(a, l, e) \ \& \ \text{state}(\gamma, t) \models \text{food_location}(l) \right]$$


7 Implications for Other Types of Society

In the previous sections, the notion of a mapping between a multi-agent and single agent conceptualisation of processes in the world has been illustrated for an example ant society. In this section, it will be shown that this example is in fact paradigmatic for a whole class of examples, including examples from human society. The general idea is as follows. Suppose certain processes are modelled as a multi-agent process, where all agents interact with a given part of the world by observing it and making changes in it. At any point in time, the state of this part of the world is the result of contributions of multiple agents. Moreover, this state affects the behaviour of the agents: their behaviour depends on this part of the world. So, at an abstract level this describes what happens in the ant society case. Examples of this beyond ant societies are:

- the blackboard in a blackboard system,
- the collection of past cases that can be considered for jurisprudence by lawyers,
- a common database in an organisation that is used to maintain and retrieve common client information,
- a common knowledge repository in the context of knowledge management
- a common guidelines document for a certain department.
- a common automated advice system for a certain department

Let us elaborate a bit on the last examples by discussing a case we encountered a few years ago within a banking organisation. A given department is responsible for advising clients for certain products P, depending on the background context C of the client. The management of the department notices that an advice is often depending on the specific advisor, and considers this as less desirable. It wants to put an effort to show as a department more 'one face' to the outside world of clients. The management considers the possibility to automate the tasks by means of a decision support system that for each context more or less forces advisors to one solution. Now, in the context of the current paper, assume that in such a situation all pieces of advice given by all employees within the department are stored in a case database DB, in the form of tuples

$$\langle C, P, E, t \rangle$$

with C a context representation, P the advised product for that context, E the employee giving this advice and t the time when the advice was given. The daily practice then should be as follows. For any employee, having a client with context C, the employee inspects the database and retrieves all previous cases for the same context:

$$\text{CASES}(C) = \{ \langle P, E, t \rangle \mid \langle C, P, E, t \rangle \in \text{DB} \}$$

If this set is empty, the employee just gives her own advice and adds this to the database. When this set is not empty, on the basis of this set it is determined what advice should be given. This can be done in a number of ways:

- take the advice P that has most frequently been given in the past: for context C take P such that
$$\max \{ \# \{ \langle P', t \rangle \mid \exists E \langle P', E, t \rangle \in \text{CASES}(C) \} \mid P' \text{ product} \} = \# \{ \langle P, t \rangle \mid \exists E \langle P, E, t \rangle \in \text{CASES}(C) \}$$
- take the advice P that has most votes of employees to be used by them: for context C take P such that
$$\max \{ \# \{ E \mid \exists t \langle P', E, t \rangle \in \text{CASES}(C) \} \mid P' \text{ product} \} = \# \{ E \mid \exists t \langle P, E, t \rangle \in \text{CASES}(C) \}$$
- take the advice P that has most frequently been given in the past, but apply weight factors according to a decay rate that values a more recent advice higher than an earlier advice: for context C and for the current time point t, take P such that
$$\max \{ \sum_i r^{-t_i} \# \{ \langle P', t' \rangle \mid \exists E \langle P', E, t' \rangle \in \text{CASES}(C) \} \mid P' \text{ product} \} = \sum_i r^{-t_i} \# \{ \langle P, t' \rangle \mid \exists E \langle P, E, t' \rangle \in \text{CASES}(C) \}$$

Notice that this last approach is, even in structure, quite similar to the ants case. The edges correspond to the pairs $\langle C, P \rangle$, and crossing an edge corresponds to giving such an advice. But also for the other approaches considered the similarity exists. By accumulating the past actions of the multiple agents, a structure is created so that the multi-agent system behaves more with one face to the external world of clients. Therefore, the clients can more easily consider the department as one single agent. The formal interpretation mapping as described in this paper formalises this situation.



8 Discussion

This paper addresses the question to what extent a process involving multiple agents that shows some form of collective intelligence can be interpreted as single agent behaviour. The question is answered by formal analysis. It is shown for an example process how it can be conceptualised and formalised in two different manners: from a single agent (or cognitive) and from a multi-agent (or social) perspective. Moreover, it is shown how a basic ontological mapping can be formally defined between the two formalisations, and how this mapping can be extended to a mapping of dynamic properties. Thus it is shown how the collective behaviour can be interpreted in a formal manner as single agent behaviour. For example, the fact that food is taken from the source to the nest can be explained by a sequence of actions of one agent, based on its beliefs. Although the case study addressed the simple example of an ant colony, it was shown that the presented interpretation approach can be applied to human societies as well. For example, the various processes going on within the department of an organisation can be explained as the behaviour of a single agent.

Having a mapping as described above allows one to explain collective or social behaviour in terms of single agent concepts, in the following manner. Behaviour often is explained by considering the basic underlying causal relations or mechanisms. The mapping (and its formalisation) allows one to replace an explanation of behaviour in terms of basic mechanisms involving frequent interactions of the multiple agents (with each other and/or with the external world), by an explanation that leaves out these interactions and bases itself directly on mental states of the single agent conceptualisation. This explanation is often simpler, more abstract, better understandable, and perhaps more elegant than the more complicated explanation based on the interactions. This is made possible by introducing a new ontology for states involved. For example, 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 reconceptualised as, for example, ‘it is believed that d is a relevant path’. Likewise, for the banking example, state properties such as ‘advice P has received the highest number of votes’ can be reconceptualised as ‘it is believed that P is an appropriate advice’.

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 (1996, 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 (1998), 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.
(Clark and Chalmers 1998), Section 3.

In (Jonker et al. 2002) it is explained in some detail 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 genericity 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; finally, the same concepts have a wider scope of application, thus obtaining unification.

Also by Dennett it is claimed that the use of a different ontology for the same world facts can be beneficial. In (Dennett 1991b), he puts forward the *intentional stance*, a perspective that allows one to describe certain physical phenomena in terms of mental concepts such as desires and intentions, in order to obtain more understandable explanations:

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 1991b), p. 42.

In this context, the perspective taken in the current paper can be viewed as an extension of the intentional stance, where mental concepts are ascribed not only to single agents, but also to processes that can be conceptualised as groups of agents. A difference with (Dennett 1991b) is that the only types of mental states addressed in the current paper are beliefs. Nevertheless, we expect that our approach can be extended in order to ascribe other mental states (such as desires and intentions) to multi-agent societies as well. Further research will have to confirm this.

Given the perspective of the intentional stance, the question might come up whether the behavioural description of the resulting ‘super-agent’ will not become just as complex as that of the initial multi-agent system. Two answers may be given to this question. First, the ants case study addressed in this paper has shown that there are at least a number of concepts in the multi-agent description that can be left out in the single agent description. In particular, such concepts are the creation and the observation of the ‘shared extended mental state’ (state m2 in Figure 3); also see Table 3 where for some concepts in the right

column there is no counterpart in the left column. Moreover, even if the single agent description is still rather complex, this does not have to be a problem. Within cognitive science, many approaches exist to handle complexity of an agent's mental processes by imposing structure on it, see, e.g., (Fodor 1983). In this view, the collective behaviour of a group of agents may be seen as single agent behaviour that consists of a number of sub-processes.

In Section 4.3, it was mentioned that the mapping from multi-agent to single agent conceptualisation is unidirectional, not bidirectional. The main reason for this was that a number of the 'collective' concepts did not have an 'individual' counterpart. However, in the literature on philosophy of mind, several authors show that in some cases it might also be beneficial to explain an *individual* mental process as a *collective* process, see e.g. (Dennett 1991a). Thus, it might be useful to explore more possibilities to obtain a mapping in the opposite direction. In future work, these possibilities will be investigated in more detail.

Other future research will further analyse the interpretation mapping in the context of logic: the notion of an interpretation of one (formal) logical theory T in another logical theory T' has a formal definition in logic. It is an interesting question whether it can be proven logically that the conditions of this definition are fulfilled for the mapping defined in this paper. For example, a question is whether it can be proven that:

$$T \Vdash \alpha \Rightarrow T' \dashv \vdash \varphi(\alpha)$$

for all formulae α , where T is a logical theory of single agent behaviour and T' a theory of multi-agent behaviour. More specifically, suppose that a global property B is implied by a number of local properties A_1, \dots, A_n , according to the following relation:

$$A_1 \& \dots \& A_n \Rightarrow B$$

Given this implication, the question to explore would be whether there is a similar relation available between the mapped properties, i.e., whether the following implication:

$$\varphi(A_1) \& \dots \& \varphi(A_n) \Rightarrow \varphi(B)$$

holds as well.



Acknowledgements

The authors are grateful to Catholijn Jonker for many valuable discussions on the topic of the paper, to Lourens van der Meij for his contribution to the development of the software environment, to Martijn Schut for his contribution to parts of the simulations, and to the anonymous referees for their comments on an earlier version of this paper.



References

1. BONABEAU, E., Dorigo, M., and Theraulaz, G. (1999) *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York.
2. 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. et al. (eds.), *Proceedings of the Third German Conference on Multi-Agent System Technologies, MATES'05*. Lecture Notes in AI, vol. 3550. Springer Verlag, pp. 165-178.
3. CLARK, A. (1997) *Being There: Putting Brain, Body and World Together Again*. MIT Press.
4. CLARK, A. and Chalmers, D. (1998) The Extended Mind. In: *Analysis*, vol. 58, pp. 7-19.
5. 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, pp. 176-186.
6. DENNETT, D.C. (1991a) *Consciousness Explained*, Little, Brown: Boston, Massachusetts.
 7. DENNETT, D.C. (1991b) Real Patterns. *The Journal of Philosophy*, vol. 88, pp. 27-51.
 8. DENNETT, D.C. (1996) *Kinds of Mind: Towards an Understanding of Consciousness*, New York: Basic Books.
 9. DROGOUL, A., Corbara, B., and Fresneau, D. (1995) MANTA: New experimental results on the emergence of (artificial) ant societies. In: Gilbert N. and Conte R. (eds.), *Artificial Societies: the computer simulation of social life*, UCL Press.
 10. FODOR, J.A. (1983) *The Modularity of Mind*, Bradford Books, MIT Press: Cambridge, Massachusetts.
 11. JONKER, C.M., Treur, J., and Wijngaards, W.C.A. (2002) Reductionist and Antireductionist Perspectives on Dynamics. *Philosophical Psychology Journal*, vol. 15, pp. 381-409.
 12. 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), pp. 191-210.
 13. KIM, J. (1996) *Philosophy of Mind*. Westview Press.
 14. KIRSH, D. and Maglio, P. (1994) On distinguishing epistemic from pragmatic action. *Cognitive Science*, vol. 18, pp. 513-49.
 15. MENARY, R. (ed.). (2006) *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, to appear.

Appendix A - The Multi-Agent Simulation Model

LP1 (Initialisation of Pheromones)

“At the start of the simulation, at all locations there are 0 pheromones.”

start \leftrightarrow pheromones_at(E1, 0.0) and pheromones_at(E2, 0.0) and pheromones_at(E3, 0.0) and pheromones_at(E4, 0.0) and pheromones_at(E5, 0.0) and pheromones_at(E6, 0.0) and pheromones_at(E7, 0.0) and pheromones_at(E8, 0.0) and pheromones_at(E9, 0.0) and pheromones_at(E10, 0.0)

LP2 (Initialisation of Ants)

“At the start of the simulation, all ants are at location A.”

start \leftrightarrow is_at_location_from(ant1, A, init) and is_at_location_from(ant2, A, init) and is_at_location_from(ant3, A, init)

LP3 (Initialisation of World)

“These two properties model the ant world. The first property expresses which locations are connected to each other, and via which edges they are connected. The second property expresses for each location how many neighbours it has.”

start \leftrightarrow connected_to_via(A, B, I1) and ... and connected_to_via(D, H, I10)

start \leftrightarrow neighbours(A, 2) and ... and neighbours(H, 3)

LP4 (Initialisation of Attractive Directions)

“This property expresses for each ant and each location, which edge is most attractive for the ant at if it arrives at that location. This criterion can be used in case an ant arrives at a location where there are two edges with an equal amount of pheromones.”

start \leftrightarrow attractive_direction_at(ant1, A, E1) and ... and attractive_direction_at(ant3, E, E5)

LP5 (Selection of Edge)

“These properties model the edge selection mechanism of the ants. For example, the first property expresses that, when an ant observes that it is at location A, and both edges connected to location A have the same number of pheromones, then the ant goes to its attractive direction.”

observes(a, is_at_location_from(A, e0)) and attractive_direction_at(a, A, e1) and connected_to_via(A, I1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(A, I2, e2) and observes(a, pheromones_at(e2, i2)) and $e1 \neq e2$ and $i1 = i2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, A, I1))

observes(a, is_at_location_from(A, e0)) and connected_to_via(A, I1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(A, I2, e2) and observes(a, pheromones_at(e2, i2)) and $i1 > i2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, A, I1))

observes(a, is_at_location_from(F, e0)) and connected_to_via(F, I1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(F, I2, e2) and observes(a, pheromones_at(e2, i2)) and $i1 > i2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, F, I1))

observes(a, is_at_location_from(I, e0)) and neighbours(I, 2) and connected_to_via(I, I1, e1) and $e0 \neq e1$ and $I \neq A$ and $I \neq F \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, I, I1))

observes(a, is_at_location_from(I, e0)) and attractive_direction_at(a, I, e1) and neighbours(I, 3) and connected_to_via(I, I1, e1) and observes(a, pheromones_at(e1, 0.0)) and connected_to_via(I, I2, e2) and observes(a, pheromones_at(e2, 0.0)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, I, I1))

observes(a, is_at_location_from(I, e0)) and neighbours(I, 3) and connected_to_via(I, I1, e1) and observes(a, pheromones_at(e1, i1)) and connected_to_via(I, I2, e2) and observes(a, pheromones_at(e2, i2)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2$ and $i1 > i2 \leftrightarrow$ to_be_performed(a, go_to_edge_from_to(e1, I1))

LP6 (Arrival at Edge)

“If an ant goes to an edge e from a location I to a location I1, then later the ant will be at this edge e.”

to_be_performed(a, go_to_edge_from_to(e, I, I1)) \leftrightarrow is_at_edge_from_to(a, e, I, I1)

LP7 (Observation of Edge)

“If an ant is at a certain edge e, going from a location I to a location I1, then it will observe this.”

is_at_edge_from_to(a, e, I, I1) \leftrightarrow observes(a, is_at_edge_from_to(e, I, I1))

LP8 (Movement to Location)

“If an ant observes that it is at an edge e from a location I to a location I1, then it will go to location I1.”

observes(a, is_at_edge_from_to(e, I, I1)) \leftrightarrow to_be_performed(a, go_to_location_from(I1, e))

LP9 (Dropping of Pheromones)

“If an ant observes that it is at an edge e from a location l to a location $l1$, then it will drop pheromones at this edge e .”

$is_at_edge_from_to(e, l, l1) \leftrightarrow to_be_performed(a, drop_pheromones_at_edge_from(e, l))$

LP10 (Arrival at Location)

“If an ant goes to a location l from an edge e , then later it will be at this location l .”

$to_be_performed(a, go_to_location_from(l, e)) \leftrightarrow is_at_location_from(a, l, e)$

LP11 (Observation of Location)

“If an ant is at a certain location l , then it will observe this.”

$is_at_location_from(a, l, e) \leftrightarrow observes(a, is_at_location_from(l, e))$

LP12 (Observation of Pheromones)

“If an ant is at a certain location l , then it will observe the number of pheromones present at all edges that are connected to location l .”

$is_at_location_from(a, l, e0) \text{ and } connected_to_via(l, l1, e1) \text{ and } pheromones_at(e1, i) \leftrightarrow observes(a, pheromones_at(e1, i))$

LP13 (Increment of Pheromones)

“These properties model the increment of the number of pheromones at an edge as a result of ants dropping pheromones. For example, the first property expresses that, if an ant drops pheromones at edge e , and no other ants drop pheromones at this edge, then the new number of pheromones at e becomes $i * decay + incr$. Here, i is the old number of pheromones, $decay$ is the decay factor, and $incr$ is the amount of pheromones dropped.”

$to_be_performed(a1, drop_pheromones_at_edge_from(e, l1)) \text{ and } \forall l2 \text{ not } to_be_performed(a2, drop_pheromones_at_edge_from(e, l2)) \text{ and } \forall l3 \text{ not } to_be_performed(a3, drop_pheromones_at_edge_from(e, l3)) \text{ and } a1 \neq a2 \text{ and } a1 \neq a3 \text{ and } a2 \neq a3 \text{ and } pheromones_at(e, i) \leftrightarrow pheromones_at(e, i * decay + incr)$

$to_be_performed(a1, drop_pheromones_at_edge_from(e, l1)) \text{ and } to_be_performed(a2, drop_pheromones_at_edge_from(e, l2)) \text{ and } \forall l3 \text{ not } to_be_performed(a3, drop_pheromones_at_edge_from(e, l3)) \text{ and } a1 \neq a2 \text{ and } a1 \neq a3 \text{ and } a2 \neq a3 \text{ and } pheromones_at(e, i) \leftrightarrow pheromones_at(e, i * decay + incr + incr)$

$to_be_performed(a1, drop_pheromones_at_edge_from(e, l1)) \text{ and } to_be_performed(a2, drop_pheromones_at_edge_from(e, l2)) \text{ and } to_be_performed(a3, drop_pheromones_at_edge_from(e, l3)) \text{ and } a1 \neq a2 \text{ and } a1 \neq a3 \text{ and } a2 \neq a3 \text{ and } pheromones_at(e, i) \leftrightarrow pheromones_at(e, i * decay + incr + incr + incr)$

LP14 (Collecting of Food)

“If an ant observes that it is at location F (the food source), then it will pick up some food.”

$observes(a, is_at_location_from(F, e)) \leftrightarrow to_be_performed(a, pick_up_food)$

LP15 (Carrying of Food)

“If an ant picks up food, then as a result it will be carrying food.”

$to_be_performed(a, pick_up_food) \leftrightarrow is_carrying_food(a)$

LP16 (Dropping of Food)

“If an ant is carrying food, and observes that it is at location A (the nest), then the ant will drop the food.”

$observes(a, is_at_location_from(A, e)) \text{ and } is_carrying_food(a) \leftrightarrow to_be_performed(a, drop_food)$

LP17 (Persistence of Food)

“As long as an ant that is carrying food does not drop the food, it will keep on carrying it.”

$is_carrying_food(a) \text{ and } \text{not } to_be_performed(a, drop_food) \leftrightarrow is_carrying_food(a)$

LP18 (Decay of Pheromones)

“If the old amount of pheromones at an edge is i , and there is no ant dropping any pheromones at this edge, then the new amount of pheromones at e will be $i * decay$.”

$pheromones_at(e, i) \text{ and } \forall a, l \text{ not } to_be_performed(a, drop_pheromones_at_edge_from(e, l)) \leftrightarrow pheromones_at(e, i * decay)$

Appendix B - The Single Agent Simulation Model

LP1' (Initialisation of Beliefs)

"At the start of the simulation, superant A believes that all locations have relevance level 0."

start \leftrightarrow belief(S, relevance_level(E1, 0.0)) and belief(S, relevance_level(E2, 0.0)) and belief(S, relevance_level(E3, 0.0)) and belief(S, relevance_level(E4, 0.0)) and belief(S, relevance_level(E5, 0.0)) and belief(S, relevance_level(E6, 0.0)) and belief(S, relevance_level(E7, 0.0)) and belief(S, relevance_level(E8, 0.0)) and belief(S, relevance_level(E9, 0.0)) and belief(S, relevance_level(E10, 0.0))

LP2' (Initialisation of Paws)

"At the start of the simulation, S has all of its paws at location A."

start \leftrightarrow has_paw_at_location_from(S, paw1, A, init) and has_paw_at_location_from(S, paw2, A, init) and has_paw_at_location_from(S, paw3, A, init)

LP3' (Initialisation of World)

"These two properties model the ant world. The first property expresses which locations are connected to each other, and via which edges they are connected. The second property expresses for each location how many neighbours it has."

start \leftrightarrow connected_to_via(A, B, I1) and ... and connected_to_via(D, H, I10)

start \leftrightarrow neighbours(A, 2) and ... and neighbours(H, 3)

LP4' (Initialisation of Attractive Directions)

"This property expresses for each paw and each location, which edge is most attractive for the paw at if it arrives at that location. This criterion can be used in case a paw arrives at a location where there are two edges with an equal amount of pheromones."

start \leftrightarrow attractive_direction_at(paw1, A, E1) and ... and attractive_direction_at(paw3, E, E5)

LP5' (Selection of Edge)

"These properties model the edge selection mechanism of superant S. For example, the first property expresses that, when S observes that it has a paw p at location A, and S believes that the relevance level of both edges connected to location A is equal, then S will move its paw to its attractive direction."

observes(S, has_paw_at_location_from(p, A, e0)) and connected_to_via(p, A, e1) and connected_to_via(A, I1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(A, I2, e2) and belief(S, relevance_level(e2, i2)) and $e1 \neq e2$ and $i1 = i2$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, A, I1))

observes(S, has_paw_at_location_from(p, A, e0)) and connected_to_via(A, I1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(A, I2, e2) and belief(S, relevance_level(e2, i2)) and $i1 > i2$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, A, I1))

observes(S, has_paw_at_location_from(p, F, e0)) and connected_to_via(F, I1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(F, I2, e2) and belief(S, relevance_level(e2, i2)) and $i1 > i2$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, F, I1))

observes(S, has_paw_at_location_from(p, I, e0)) and neighbours(I, 2) and connected_to_via(I, I1, e1) and $e0 \neq e1$ and $I \neq A$ and $I \neq F$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, I, I1))

observes(S, has_paw_at_location_from(p, I, e0)) and attractive_direction_at(a, I, e1) and neighbours(I, 3) and connected_to_via(I, I1, e1) and belief(S, relevance_level(e1, 0.0)) and connected_to_via(I, I2, e2) and belief(S, relevance_level(e2, 0.0)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, I, I1))

observes(S, has_paw_at_location_from(p, I, e0)) and neighbours(I, 3) and connected_to_via(I, I1, e1) and belief(S, relevance_level(e1, i1)) and connected_to_via(I, I2, e2) and belief(S, relevance_level(e2, i2)) and $e0 \neq e1$ and $e0 \neq e2$ and $e1 \neq e2$ and $i1 > i2$ \leftrightarrow to_be_performed(S, move_paw_to_edge_from_to(p, e1, I1))

LP6' (Paw Arrival at Edge)

"If S moves its paw p to an edge e from a location I to a location I1, then later this paw will be at this edge e."

to_be_performed(S, move_paw_to_edge_from_to(p, e, I, I1)) \leftrightarrow has_paw_at_edge_from_to(S, p, e, I, I1)

LP7' (Paw Observation at Edge)

"If S has a paw at a certain edge e, going from a location I to a location I1, then it will observe this."

has_paw_at_edge_from_to(S, p, e, I, I1) \leftrightarrow observes(S, has_paw_at_edge_from_to(p, e, I, I1))

LP8' (Paw Movement to Location)

'If S observes that it has a paw p at an edge e from a location l to a location l1, then it will move this paw to location l1.'

observes(S, has_paw_at_edge_from_to(p, e, l, l1)) \leftrightarrow to_be_performed(S, move_paw_to_location_from(p, l1, e))

LP9' (Paw Arrival at Location)

'If S moves its paw p to a location l from an edge e, then later this paw will be at this location l.'

to_be_performed(S, move_paw_to_location_from(p, l, e)) \leftrightarrow has_paw_at_location_from(S, p, l, e)

LP10' (Paw Observation at Location)

'If S has a paw p at a certain location l, then it will observe this.'

has_paw_at_location_from(S, p, l, e) \leftrightarrow observes(S, has_paw_at_location_from(p, l, e))

LP11' (Increment of Belief)

'These properties model the increment of S' belief of the relevance level of an edge as a result of the presence of its paws there. For example, the first property expresses that, if S has exactly one paw at edge e, then the new number of pheromones at e becomes $i * \text{decay} + \text{incr}$. Here, i is the old relevance level, decay is the decay factor, and incr is the increment value of the belief.'

observes(S, has_paw_at_edge_from_to(p1, e, l, l1)) and $\forall l2$ not observes(S, has_paw_at_edge_from_to(p2, e, l, l2)) and $\forall l3$ not observes(S, has_paw_at_edge_from_to(p3, e, l, l3)) and $p1 \neq p2$ and $p1 \neq p3$ and $p2 \neq p3$ and belief(S, relevance_level(e, i)) \leftrightarrow belief(S, relevance_level(e, $i * \text{decay} + \text{incr}$))

observes(S, has_paw_at_edge_from_to(p1, e, l, l1)) and observes(S, has_paw_at_edge_from_to(p2, e, l, l2)) and $\forall l3$ not observes(S, has_paw_at_edge_from_to(p3, e, l, l3)) and $p1 \neq p2$ and $p1 \neq p3$ and $p2 \neq p3$ and belief(S, relevance_level(e, i)) \leftrightarrow belief(S, relevance_level(e, $i * \text{decay} + \text{incr} + \text{incr}$))

observes(S, has_paw_at_edge_from_to(p1, e, l, l1)) and observes(S, has_paw_at_edge_from_to(p2, e, l, l2)) and observes(S, has_paw_at_edge_from_to(p3, e, l, l3)) and $p1 \neq p2$ and $p1 \neq p3$ and $p2 \neq p3$ and belief(S, relevance_level(e, i)) \leftrightarrow belief(S, relevance_level(e, $i * \text{decay} + \text{incr} + \text{incr} + \text{incr}$))

LP12' (Collecting of Food)

'If S observes that it has a paw p at location F (the food source), then it will pick up some food with that paw.'

observes(S, has_paw_at_location_from(p, F, e)) \leftrightarrow to_be_performed(S, pick_up_food_with_paw(p))

LP13' (Carrying of Food)

'If S picks up food with a paw p, then as a result it will be carrying food with that paw.'

to_be_performed(S, pick_up_food_with_paw(p)) \leftrightarrow is_carrying_food_with_paw(S, p)

LP14' (Dropping of Food)

'If S is carrying food with a paw p, and observes that this paw is at location A (the nest), then S will drop the food with that paw.'

observes(S, has_paw_at_location_from(p, A, e)) and is_carrying_food_with_paw(S, p) \leftrightarrow to_be_performed(S, drop_food_with_paw(p))

LP15' (Persistence of Food)

'As long as a paw that is carrying food does not drop the food, it will keep on carrying it.'

is_carrying_food_with_paw(S, p) and not to_be_performed(S, drop_food_with_paw(p)) \leftrightarrow is_carrying_food_with_paw(S, p)

LP16' (Decay of Belief)

'If S believes that the relevance level of an edge is i, and S does not observe any of its paws at this edge, then it will belief that the new relevance level of e is $i * \text{decay}$.'

belief(S, relevance_level(e, i)) and $\forall p, l, l1$ not observes(S, has_paw_at_edge_from_to(p, e, l, l1)) \leftrightarrow belief(S, relevance_level(e, $i * \text{decay}$))