On the Complexity Monotonicity Thesis for Environment, Behaviour and Cognition

Tibor Bosse, Alexei Sharpanskykh, and Jan Treur

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

Abstract. Development of more complex cognitive systems during evolution is sometimes viewed in relation to environmental complexity. In more detail, growth of complexity during evolution can be considered for the dynamics of externally observable behaviour of agents, for their internal cognitive systems, and for the environment. This paper explores temporal complexity for these three aspects, and their mutual dependencies. A number of example scenarios have been formalised in a declarative temporal language, and the complexity of the structure of the different formalisations was measured. Thus, some empirical evidence was provided for the thesis that for more complex environments, more complex behaviour and more complex mental capabilities are needed.

1 Introduction

Behaviour of agents (both living organisms and artificial (software or hardware) agents) can occur in different types and complexities, varying from very simple behaviour to more sophisticated forms. Depending on the complexity of the externally observable behaviour, the internal mental representations and capabilities required to generate the behaviour also show a large variety in complexity. From an evolutionary viewpoint, for example, Wilson [16], p. 187 and Darwin [3], p. 163 point out how the development of behaviour relates to the development of more complex cognitive capabilities. Godfrey-Smith [4], p. 3 assumes a relationship between the complexity of the environment and the development of mental representations and capabilities. He formulates the main theme of his book in condensed form as follows: 'The function of cognition is to enable the agent to deal with environmental complexity' (the *Environmental Complexity Thesis*). In this paper, this thesis is refined as follows:

- the more complex the environment, the more sophisticated is the behaviour required to deal with this environment,
- the more sophisticated the behaviour, the more complex are the mental representations and capabilities needed

This refined thesis will be called the *Complexity Monotonicity Thesis*. The idea is that to deal with the physical environment, the evolution process has generated and still generates a variety of organisms that show new forms of behaviour. These new forms of behaviour are the result of new architectures of organisms, including cognitive systems with mental representations and capabilities of various degrees of complexity. The occurrence of such more complex architectures for organisms and the induced more complex behaviour itself increases the complexity of the environment during the evolution process. New organisms that have to deal with the behaviour of such already occurring organisms live in a more complex environment, and therefore need more complex behaviour to deal with this environment, (to be) realised by an architecture with again more complex mental capabilities. In particular, more complex environments often ask for taking into account more complex histories, which requires more complex internal cognitive representations and dynamics, by which more complex behaviour is generated.

This perspective generates a number of questions. First, how can the Complexity Monotonicity Thesis be formalised, and in particular how can the 'more complex' relation be formalised for (1) the environment, (2) externally observable agent behaviour and (3) internal cognitive dynamics? Second, connecting the three items, how to formalise (a) when does a behaviour fit an environment: which types of externally observable behaviours are sufficient to cope with which types of environments, and (b) when does a cognitive system generate a certain behaviour: which types of internal cognitive dynamics are sufficient to generate which types of externally observable agent behaviour?

In this paper these questions are addressed from a dynamics perspective, and formalised by a declarative temporal logical approach. Four cases of an environment, suitable behaviour and realising cognitive system are described, with an increasing complexity over the cases. Next, for each case, complexity of the dynamics of environment, externally observable agent behaviour and internal cognitive system are formalised in terms of structure of the formalised temporal specifications describing them, thus answering (1) to (3). Moreover, (a) and (b) are addressed by establishing formalised logical (entailment) relations between the respective temporal specifications. By comparing the four cases with respect to complexity, the Complexity Monotonicity Thesis is tested.

2 Evolutionary Perspective

The environment imposes certain requirements that an agent's behaviour needs to satisfy; these requirements change due to changing environmental circumstances. The general pattern is as follows. Suppose a certain goal G for an agent (e.g., sufficient food uptake over time) is reached under certain environmental conditions ES1 (Environmental Specification 1), due to its Behavioural Specification BS1, realised by its internal (architecture) CS1 (Cognitive Specification 1). In other words, the behavioural properties BS1 are sufficient to guarantee G under environmental conditions ES1, formally ES1 & BS1 $\Rightarrow$ G, and the internal dynamics CS1 are sufficient to guarantee BS1, formally CS1 $\Rightarrow$ BS1. In other environmental circumstances, described by environmental specification ES2 (for example, more complex) the old circumstances ES1 may no longer hold, so that the goal G may no longer be reached by behavioural properties BS1. An environmental change from ES1 to ES2 may entail that behaviour BS1 becomes insufficient. It has to be replaced by new behavioural properties BS2 (also more complex) which express how under environment ES2 goal G can be achieved, i.e., ES2 & BS2 $\Rightarrow$ G.

Thus, a population is challenged to realise such behaviour BS2 by changing its internal architecture and its dynamics, and as a consequence fulfill goal G again. This challenge expresses a redesign problem: the given architecture of the agent as described by CS1 (which entails the old behavioural specification BS1) is insufficient to entail the new behavioural requirements BS2 imposed by the new environmental circumstances ES2; the evolution process has to redesign the architecture into one with internal dynamics described by some CS2 (also more complex), with CS2 $\Rightarrow$ BS2, to realise the new requirements on behaviour.

Based on these ideas, the Complexity Monotonicity Thesis can be formalised in the following manner. Suppose $< E_1, B_1, C_1 >$ and $< E_2, B_2, C_2 >$ are triples of environment, behaviour and cognitive system, respectively, such that the behaviours B_i are adequate for the respective environment E_i and realised by the cognitive system C_i. Then the Complexity Monotonicity Thesis states that

$$E_1 \leq_c E_2 \;\Rightarrow\; B_1 \leq_c B_2 \quad \& \quad B_1 \leq_c B_2 \;\Rightarrow\; C_1 \leq_c C_2$$

Here $\leq_c$ is a partial ordering in complexity, where $X \leq_c Y$ indicates that Y is more complex than X. A special case is when the complexity ordering is assumed to be a total ordering where for every two elements X, Y either $X \leq_c Y$ or $Y \leq_c X$ (i.e., they are comparable), and when some complexity measure cm is available, assigning degrees of complexity to environments, behaviours and cognitive systems, such that

$$X \leq_c Y \Leftrightarrow cm(X) \leq cm(Y)$$

where $\leq$ is the standard ordering relation on (real or natural) numbers. In this case the Complexity Monotonicity Thesis can be reformulated as

$$cm(E_1) \leq cm(E_2) \;\Rightarrow\; cm(B_1) \leq cm(B_2) \quad \&$$
$$cm(B_1) \leq cm(B_2) \;\Rightarrow\; cm(C_1) \leq cm(C_2)$$

The Temporal Complexity Monotonicity Thesis can be used to explain increase of complexity during evolution in the following manner. Make the following assumption on Addition of Environmental Complexity by Adaptation, as described above:

- adaptation of a species to an environment adds complexity to this environment

Suppose an initial environment is described by ES0, and the adapted species by BS0. Then this transforms ES0 into a more complex environmental description ES1. Based on ES1, the adapted species will have description BS1. As ES1 is more complex than ES0, by the Complexity Monotonicity Thesis it follows that this BS1 is more complex than BS0: ES0 $\leq$ ES1 $\Rightarrow$ BS0 $\leq$ BS1. Therefore BS1 again adds complexity to the environment, leading to ES2, which is more complex than ES1, et cetera[1]:

$$\begin{tikzcd}
ES0 \arrow[d] & \leq & ES1 \arrow[d] & \leq & ES2 \arrow[d] & \dots \\
BS0 \arrow[urr] & \leq & BS1 \arrow[urr] & \leq & BS2 & \dots
\end{tikzcd}$$

This argument shows that the increase of complexity during evolution can be related to and explained by two assumptions: the Complexity Monotonicity Thesis,

[1] Note that this argument can also be applied to multiple species at the same time, i.e., species A increases the complexity of the environment, which causes another species B to adapt to this more complex environment.

and the Addition of Environmental Complexity by Adaptation assumption. This paper focuses on the former assumption.

3 Variations in Behaviour and Environment

To evaluate the approach put forward, a number of cases of increasing complexity are analysed, starting from very simple *stimulus-response behaviour* solely depending on stimuli the agent gets as input at a given point in time. This can be described by a very simple temporal structure: direct associations between the input state at one time point and the (behavioural) output state at a next time point. A next class of behaviours, with slightly higher complexity, analysed is *delayed response behaviour*: behaviour that not only depends on the current stimuli, but also may depend on input of the agent in the past. This pattern of behaviour cannot be described by direct functional associations between one input state and one output state; it increases temporal complexity compared to stimulus-response behaviour. For this case, the description relating input states and output states necessarily needs a reference to inputs received in the past. Viewed from an internal perspective, to describe mental capabilities generating such a behaviour, often it is assumed that it involves a memory in the form of an internal model of the world state. Elements of this world state model mediate between the agent's input and output states.

Other types of behaviour go beyond the types of reactive behaviour sketched above. For example, behaviour that depends in a more indirect manner on the agent's input in the present or in the past. Observed from the outside, this behaviour seems to come from within the agent itself, since no direct relation to current inputs is recognised. It may suggest that the agent is motivated by itself or acts in a goal-directed manner. For a study in *goal-directed behaviour* and foraging, see, for example, [5]. Goal-directed behaviour to search for invisible food is a next case of behaviour analysed. In this case the temporal description of the externally observable behavioural dynamics may become still more complex, as it has to take into account more complex temporal relations to (more) events in the past, such as the positions already visited during a search process. Also the internal dynamics may become more complex. To describe mental capabilities generating such a type of behaviour from an internal perspective, a mental state property *goal* can be used. A goal may depend on a history of inputs. Finally, a fourth class of behaviour analysed, which also goes beyond reactive behaviour, is *learning behaviour* (e.g., conditioning). In this case, depending on its history comprising a (possibly large) number of events, the agent's externally observable behaviour is tuned. As this history of events may relate to several time points during the learning process, this again adds temporal complexity to the specifications of the behaviour and of the internal dynamics.

To analyse these four different types of behaviour in more detail, four cases of a food supplying environment are considered in which suitable food gathering behaviours are needed. These cases are chosen in such a way that they correspond to the types of behaviour mentioned above. For example, in case 1 it is expected that stimulus-response behaviour is sufficient to cope with the environment, whilst in case 2, 3 and 4, respectively, delayed response behaviour, goal-directed behaviour, and learning behaviour is needed). The basic setup is inspired by experimental literature in animal behaviour such as [6], [14], [15]. The world consists of a number of positions which have distances to each other. The agent can walk over these positions. Time is partitioned in fixed periods (days) of a duration of d time units (hours). Every day the environment generates food at certain positions, but this food may or may not be visible, accessible and persistent at given points in time. The four different types of environment with increasing temporal complexity considered are:

(1) Food is always visible and accessible. It persists until it is taken.
(2) Food is visible at least at one point in time and accessible at least at one later time point. It persists until it is taken.
(3) Food either is visible at least at one point in time and accessible at least at one later time point, or it is invisible and accessible the whole day. It persists until it is taken.
(4) One of the following cases holds:
    a) Food is visible at least at one point in time and accessible at least at one later time point. It persists until it is taken.
    b) Food is invisible and accessible the whole day. It persists until it is taken.
    c) Food pieces can disappear, and new pieces can appear, possibly at different positions. For every position where food appears, there are at least three different pieces in one day. Each piece that is present is visible. Each position is accessible at least after the second food piece disappeared.

Note that there is an accumulating effect in the increase of complexity of these types of environment. For example, the behaviour of environment (3) is described as the disjunction of the behaviour of environment (2) and another type of behaviour. For this reason, it is expected that agents that survive in environment n will also survive in environment n-1.

4 Modelling Approach

To express formal specifications for environmental, behavioural and cognitive dynamics for agents, the Temporal Trace Language (TTL, see [2]) is used. This language is a variant of order-sorted predicate logic. In dynamic property expressions, TTL allows explicit references to time points and traces. If a is a state property, then, for example state(γ, t, input(agent)) |= a denotes that this state property holds in trace γ at time point t in the input state of the agent. Here, a *trace* (or trajectory) is defined as a time-indexed sequence of states, where time points can be expressed, for example, by real or integer values. If these states are input states, such a trace is called an *input trace*. Similarly for an *output trace*. Moreover, an *input-output correlation* is defined as a binary relation C : Input_traces x Output_traces between the set of possible input traces and the set of possible output traces.

In the following sections, the four variations in behaviour and environment as introduced above are investigated in more detail. For formalising dynamic properties in TTL that will be used to specify these cases, the following state properties are used:

at(o, p)	object o is at position p
visible(sp)	an object occurring in the state property sp is visible (e.g. as it is not covered by a large object)
accessible(p)	position p is accessible (e.g. because there is no enemy at the position)
distance(p1, p2, i)	the distance between positions p1 and p2 is i
max_dist	a constant indicating the maximum distance the agent can travel in one step
observed(sp)	the agent observes state property sp
performing_action(a)	the agent performs action a

For example, a property that describes stimulus-response behaviour of an agent that goes to food, observed in the past can be expressed and formalised as follows:

At any point in time t,
if the agent observes itself at position p
and it observes an amount of food x at position p'
and position p' is accessible
then at the next time point after t the agent will go to position p'

Formalisation:
$\forall t\ \forall x\ \forall p\ \forall p'$
[state(γ, t, input(agent)) |= observed(at(agent, p)) ∧ observed(at(food(x), p')) ∧
observed(accessible(p')) ⇒ state(γ, t+1, output(agent)) |= performing_action(goto(p'))]

5 Behavioural Cases

Using the introduced approach to formalise dynamic properties, the four variations in behaviour and environment are addressed in this section: stimulus-response, delayed-response, goal-directed, and learning behaviour.

5.1 Stimulus-Response Behaviour

As a first, most simple type of behaviour, stimulus-response behaviour is analysed in more detail. For this and the following cases of behaviour the following basis properties EP1-EP5 are used to describe the behaviour of the environment. They are specified both in a structured semi-formal temporal language, and in the formal temporal language TTL. Additionally, for every case specific properties of the environment will be specified.

Environmental properties

EP1 Sufficient food within reach
At the beginning of every day n (d is the duration of a day), the agent is positioned at a position p, and a sufficient amount x of food (c is the minimum) is provided at some position p' within reachable distance from p.
∀n ∃p ∃p' ∃x ∃i x>c & i≤max_dist &
state(γ, n*d, environment) |= at(agent, p) ∧ at(food(x), p') ∧ distance(p, p', i)

EP2 Complete observability
If the agent is at position p, and a(p, p') is a visible state property involving p and a position p' within reachable distance, then this is observed by the agent. This property is to be applied to food, distance, accessibility, agent position, and the absence of these.
∀t ∀x ∀p ∀p' ∀i
[[i≤max_dist & state(γ, t, environment) |= at(agent, p) ∧ a(p, p') ∧ visible(a(p, p')) ∧
distance(p, p', i)] ⇒ state(γ, t, input(agent)) |= observed(a(p, p'))]

EP3 Guaranteed effect of movement
At any point in time t, if the agent goes to position p, then it will be at position p.
∀t ∀p state(γ, t, output(agent)) |= performing_action(goto(p))
⇒ state(γ, t+1, environment) |= at(agent, p)

EP4 Guaranteed effect of eating
At any point in time t, if the agent takes food and the amount of food is sufficient for the agent then the agent will be well fed
$\forall t$ [[$\forall x$ state(γ, t, output(agent))|= performing_action(take(food(x))) & $x \geq c$]
$\Rightarrow$ state(γ, t+1, environment) |= agent_well_fed]

EP5 Reachability of environment
The distances between all positions p in the agent's territory are smaller than max_dist. Here, p and p' are variables over the type TERRITORY_POSITION, which is a subtype of POSITION.
$\forall t\ \forall p\ \forall p'\ \forall I$ state(γ, t, environment) |= distance(p, p', i) $\Rightarrow$ i $\leq$ max_dist

The following environmental properties hold for the stimulus-response case and some of the other cases considered.

EP6 Food persistence
Food persists until taken by the agent.
$\forall t1\ \forall t2\ \forall x\ \forall p$ [t1<t2 & state(γ, t1, environment) |= at(food(x), p) &
[$\forall t$ t1 $\leq$ t $\leq$ t2 $\Rightarrow$ state(γ, t, output(agent)) |= not(performing_action(take(food(x))))]
$\Rightarrow$ state(γ, t2, environment) |= at(food(x), p)]

EP7 Food on one position
Per day, food only appears on one position.
$\forall n\ \forall x\ \forall p\ \forall p'\ \forall t$ state(γ, n*d, environment) |= at(food(x), p) &
state(γ, t, environment) |= at(food(x), p') & n*d < t $\leq$ (n+1)*d $\Rightarrow$ p = p'

EP8 Complete accessibility
Each position is accessible for the agent (i.e., never blocked by enemies).
$\forall t\ \forall p$ state(γ, t, environment) |= accessible(p)

EP9 Complete visibility
All state properties a(p, p') that are true, are visible (which means that they will be observed by agents that are close enough, according to EP2). This property is to be applied to food, distance, accessibility, agent position, and the absence of these.
$\forall t\ \forall p\ \forall p'$ state(γ, t, environment) |= a(p, p') $\Rightarrow$ state(γ, t, environment(agent)) |= visible(a(p, p'))

Note that the property of an agent being well fed is assumed to be a state property of the environment, since it refers to the agent's body state.

For the case of stimulus-response behaviour the environment is characterised by the following conjunction ES1 of a subset of the environmental properties given above:

ES1 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & EP6 & EP7 & EP8 & EP9

Behavioural Properties
The agent's stimulus-response behaviour is characterised by the following behavioural properties.

BP1 Going to observed food
At any point in time t, if the agent observes itself at position p and it observes no food at position p and it observes that an amount of food x is present at position p' and it observes that position p' is accessible and it observes that position p' is within reachable distance then it will go to position p'.
$\forall t\ \forall x\ \forall p\ \forall p'$ [[state(γ, t, input(agent)) |= observed(at(agent, p)) $\wedge$ observed(not(at(food(x), p))) $\wedge$ observed(at(food(x), p')) $\wedge$ observed(accessible(p')) $\wedge$ observed(distance(p, p', i)) & i$\leq$max_dist]
$\Rightarrow$ state(γ, t+1, output(agent)) |= performing_action(goto(p'))]

BP2 Food uptake
At any point in time t, if the agent observes itself at position p and the agent observes food at p then it will take the food
$\forall t\ \forall x\ \forall p$ [[state(γ, t, input(agent)) |= observed(at(agent, p)) $\wedge$ observed(at(food(x), p))]
$\Rightarrow$ state(γ, t+1, output(agent))|= performing_action(take(food(x)))]

Vitality property VP
The animal gets sufficient food within any given day.
$\forall$n $\exists$t1 [n*d ≤ t1 ≤ (n+1)*d & state(γ, t1, environment) |= agent_well_fed]

Logical relations
Given the dynamic properties specified above, the *environmental* and *behavioural* specifications (in short, ES1 and BS1) for case 1 (stimulus-response behaviour) are as follows:

ES1 ≡ EP1 & EP2 & EP3 & EP4 & EP5 &EP6 & EP7 & EP8 & EP9
BS1 ≡ BP1 & BP2

Given these specifications, the question is whether they are logically related in the sense that this behaviour is adequate for this environment, i.e., whether indeed the following implication holds:

BS1 & ES1 ⇒ VP

To automatically check such implications between dynamic properties at different levels, model checking techniques can be used. To this end, first the dynamic properties should be converted from TTL format to a finite state transition format. This can be done using an automated procedure, as described in [11]. After that, for checking the implications between the converted properties, the model checker SMV is appropriate (see URL: http://www.cs.cmu.edu/~modelcheck/smv.html; see also [8]). SMV has been used to verify (and confirm) the above implication, as well as a number of other implications shown in this paper.

Concerning the relation between the specification of the *cognitive* and the *behavioural* dynamics: in this case CS1 = BS1. Thus, CS1 ⇒ BS1 also holds.

5.2 Delayed Response Behaviour

In delayed response behaviour, previous observations may have led to maintenance of some form of memory of the world state: a model or representation of the (current) world state (for short, *world state model*). This form of memory can be used at any point in time as an additional source (in addition to the direct observations). In that case, at a given time point the same input of stimuli can lead to different behavioural output, since the world state models based on observations in the past can be different. This makes that agent behaviours do not fit in the setting of an input-output correlation based on a direct functional association between (current) input states and output states. Viewed from an external viewpoint, this type of behaviour, which just like stimulus-response behaviour occurs quite often in nature, is just a bit more complex than stimulus-response behaviour, in the sense that it adds complexity to the temporal dimension by referring not only to current observations but also to observations that took place in the past.

This leads to the question what kind of complexity in the environment is coped with this kind of behaviour that is not coped with by stimulus-response behaviour. An answer on this question can be found in a type of environment with aspects which are important for the animal (e.g., food or predators), and which cannot be completely observed all the time; e.g., food or predators are sometimes hidden by other objects:

Environmental properties

For this case the environment described sometimes shows the food, but not always as in the previous case. It is characterised by the following conjunction ES2 of a subset of the environmental properties given above, extended with the properties EP10, EP11 and EP12 given below:

ES2 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & EP6 & EP7 & EP10 & EP11 & EP12

EP10 Temporary visibility of food

Per day, all food that is present is visible for at least one time point, and is accessible for at least one later time point[2].

EP11 Complete visibility of non-food

All state properties that are true, except the presence of food, are visible. Thus, this property is applied to distance, accessibility, and agent position.

EP12 Complete local observability of food

For all time points, if the agent is at the position p with food then the agent observes the food (no matter if it is visible, e.g., by smell)

Behavioural properties

Next, dynamic properties are identified that characterise the input-output correlation of delayed response behaviour, observed from an external viewpoint. Such a dynamic property has a temporal nature; it can refer to the agent's input and output in the present, the past and/or the future. In semi-formal and formal notation, for the case considered, the input-output correlation for delayed response behaviour can be characterised by:

BP3 Going to food observed in the past

At any point in time t, if the agent observes itself at position p and it observes no food at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and at some earlier point in time t1 the agent observed that an amount of food x was present at position p' and at every point in time t2 after t1 up to t, the agent did not observe that no food was present at p' then at the next time point after t the agent will go to position p'

∀t ∀x ∀i ∀p ∀p'
[[state(γ, t, input(agent)) |= observed(at(agent, p)) ∧ observed(not(at(food(x), p))) ∧
observed(accessible(p')) ∧ observed(distance(p, p', i)) & i≤max_dist] &
∃t1<t [state(γ, t1, input(agent)) |= observed(at(food(x), p')) &
∀t2 [t ≥ t2 > t1 ⇒ state(γ, t2, input(agent))|= not(observed(not(at(food(x), p'))))]]
⇒ state(γ, t+1, output(agent)) |= performing_action(goto(p'))]

Cognitive properties

Since the external characterisations of delayed response behaviour refer to the agent's input in the past, it is assumed that internally the agent maintains past observations by means of persisting internal state properties, i.e., some form of memory. These persisting state properties are sometimes called *beliefs*. For the example case, it is assumed that an internal state property b1(p) is available, with the following dynamics:

CP1 Belief formation on food presence

At any point in time t, if the agent observes that food is present at position p then internal state property b1(p) will hold (i.e., a belief that food is present at p)

[2] Formal expressions for all properties can be found in the Appendix at http://www.cs.vu.nl/~tbosse/complexity.

CP2 Belief b1 persistence
At any point in time t, if internal state property b1(p) holds and the agent does not observe the absence of food at position p then at the next time point internal state property b1(p) still holds

CP3 Going to food believed present
At any point in time t, if the agent observes itself at position p and it observes no food at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and $p \neq p'$ and internal state property b1(p') holds then the agent will go to position p'

Logical relations

ES2 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & EP6 & EP7 & EP10 & EP11 & EP12
BS2 ≡ BP2 & BP3
CS2 ≡ BP2 & CP1 & CP2 & CP3
BS2 & ES2 ⇒ VP
CS2 ⇒ BS2

5.3 Goal-Directed Behaviour

A next, more complex type of behaviour considered is goal-directed behaviour. This behaviour is able to cope with environments where visibility can be more limited than in the environments considered before.

Environmental properties

For this case the environment is characterised by the following expression ES3 based on a subset of the environmental properties given earlier, extended with property EP13, given below:

ES3 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & EP6 & EP7 & EP11 & EP12 &
(EP10 OR (EP8 & EP13))

EP13 Complete invisibility of food
Food is always invisible for the agent (e.g., always covered), unless the agent is at the same position as the food.

Behavioural properties

The agent's behaviour exploring positions in order to discover food is characterised by the following behavioural property:

BP4 Searching for food
At any point in time t, if the agent observes itself at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and it did not visit position p' yet and p' is the position closest to p which the agent did not visit and it did not observe any food at all yet then at the next time point after t the agent will go to position p'

$\forall t\ \forall p\ \forall p'$
state(γ, t, input(agent)) |= observed(at(agent, p)) ∧ observed(accessible(p')) ∧
observed(distance(p, p', i)) & i≤max_dist &
not [∃t' t'<t & state(γ, t', input(agent)) |= observed_at(agent, p')] &
∀p'' [[not [∃t' t'<t & state(γ, t', input(agent)) |= observed_at(agent, p'')]]
⇒ ∃d1 ∃d2 state(γ, t, input(agent)) |= observed(distance(p, p', d1)) ∧
observed(distance(p, p'', d2)) & d1<d2] &
not [∃t' ∃p'' ∃x t'≤t & state(γ, t', input(agent)) |= observed(at(food(x), p''))]
⇒ state(γ, t+1, output(agent)) |= performing_action(goto(p'))

Cognitive properties

To describe the internal cognitive process generating this type of behaviour, the mental state property *goal* is used. In particular, for the case addressed here, when the agent has no beliefs about the presence of food, it will generate the goal to find food. If it has this goal, it will pro-actively search for food in unexplored positions. This is characterised by the following dynamic properties:

CP4 Goal formation
At any point in time t, if the agent does not believe that food is present at any position p then it will have the goal to find food

CP5 Non-goal formation
At any point in time t, if the agent believes that food is present at position p then it will not have the goal to find food

CP6 Belief formation on visited position
At any point in time t, if the agent observes itself at position p then internal state property b2(p) will hold (i.e., the belief that it visited p)

CP7 Belief b2 persistence
At any point in time t, if internal state property b2(p) holds then at the next time point internal state property b2(p) still holds

CP8 Belief formation on distances
At any point in time t, if the agent observes that the distance between position p and p' is d then internal state property belief(p, p', d) will hold

CP9 Belief persistence on distances
At any point in time t, if internal state property belief(p, p', d) holds then at the next time point internal state property belief(p, p', d) still holds

CP10 Going to closest position
At any point in time t, if the agent observes itself at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and it has the goal to find food and it believes it did not visit p' yet and p' is the position closest to p of which the agent believes it did not visit it then at the next time point after t the agent will go to position p'

Logical relations

ES3 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & EP6 & EP7 & EP11 & EP12 &
(EP10 OR (EP8 & EP13))
BS3 ≡ BP2 & BP3 & BP4
CS3 ≡ BP2 & CP1 & CP2 & CP3 & CP4 & CP5 & CP6 & CP7 & CP8 & CP9 & CP10
BS3 & ES3 $\Rightarrow$ VP
CS3 $\Rightarrow$ BS3

5.4 Learning Behaviour

A final class of behaviour analysed is learning behaviour. In this case, depending on its history comprising a (possibly large) number of events, the agent’s externally observable behaviour is tuned to the environment’s dynamics. In the case addressed here, in contrast to the earlier cases, the environment has no guaranteed persistence of food for all positions. Instead, at certain positions food may come and go (e.g., because it is eaten by competitors). The agent has to learn that, when food often appears (and disappears) at a certain position, then this is an interesting position to be, because food may re-appear at that position (but soon disappear again).

Environmental properties

For this case the environment is characterised by the following expression ES4 based on a subset of the environmental properties given earlier, extended with property EP14, given below.

ES4 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & ((EP6 & EP7 & EP10 & EP11 & EP12)
OR (EP6 & EP7 & EP8 & EP11 & EP12 & EP13) OR (EP9 & EP14))

EP14 Food reoccurrence

Every piece of food disappears and reappears at least 2 times per day, of which at least after the second disappearance its position will be accessible.

Behavioural properties

The agent's behaviour for this case should take into account which positions show reoccurence of food. The following behavioural property characterises this.

BP5 Being at useful positions

At any point in time t, if the agent observes itself at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and for all positions p" that the agent observed food in the past, the agent later observed that the food disappeared and at some earlier point in time t1 the agent observed that food was present at position p' and after that at time point t2 before t the agent observed no food present at position p' and after that at time point t3 before t the agent again observed the presence of food at position p' and after that at a time point t4 before t the agent again observed no food present at position p' and p' is the closest reachable position for which the above four conditions hold then at the next time point after t the agent will go to position p'

∀t ∀p ∀p' ∀x
state(γ, t, input(agent)) |= observed(at(agent, p)) ∧
observed(accessible(p')) ∧ observed(distance(p, p', i)) & i≤max_dist &
∀t' ∀p'' ∀x' [t'<t & state(γ, t', input(agent)) |= observed(at(food(x'), p''))
⇒ ∃t'' t'<t''≤t &
state(γ, t'', input(agent)) |= observed(not(at(food(x'), p'')))]
& ∃t1 ∃t2 ∃t3 ∃t4 [t1<t2<t3<t4<t &
state(γ, t1, input(agent)) |= observed(at(food(x), p')) &
state(γ, t2, input(agent)) |= observed(not(at(food(x), p'))) &
state(γ, t3, input(agent)) |= observed(at(food(x), p')) &
state(γ, t4, input(agent)) |= observed(not(at(food(x), p')))]
& ∀p" [∃t1 ∃t2 ∃t3 ∃t4 [t1<t2<t3<t4 &
state(γ, t1, input(agent)) |= observed(at(food(x), p")) &
state(γ, t2, input(agent)) |= observed(not(at(food(x), p"))) &
state(γ, t3, input(agent)) |= observed(at(food(x), p")) &
state(γ, t4, input(agent)) |= observed(not(at(food(x), p")))] ⇒
∃d1 ∃d2
state(γ, t, input(agent)) |= observed(distance(p, p', d1)) ∧
observed(distance(p, p'', d2)) & d1<d2]
⇒ state(γ, t+1, output(agent)) |= performing_action(goto(p'))

Cognitive properties

The internal cognitive dynamics has to take into account longer histories of positions and food (re)appearing there. This is realised by representations that are built up for more complex world properties, in particular, not properties of single states but of histories of states of the world. For example, at a certain time point, it has to be represented that for a certain position in the past food has appeared twice and in between disappeared. The state properties b3(p, q) play the role of representations of world histories on food (re)occurrence.

CP11 Initial mental state
At the beginning of every day n, for all positions p, internal state property b3(p, 0) holds (i.e. a belief that there is no food at p)

CP12 Belief update on food presence
At any point in time t, for $q \in \{0,2\}$, if internal state property b3(p, q) holds and the agent observes food at position p then internal state property b3b(p, q+1) will hold

CP13 Belief update on food absence
At any point in time t, for $q \in \{1,3\}$, if internal state property b3(p,q) holds and the agent observes no food at position p then internal state property b3(p,q+1) will hold

CP14 Belief b3 persistence
At any point in time t, for all q, if internal state property b3(p,q) holds then at the next time point internal state property b3(p,q) still holds

CP15 Going to interesting position
At any point in time t, if the agent observes itself at position p and it observes that position p' is accessible and it observes that position p' is within reachable distance and it has the goal to find food and p' is the position closest to p of which the agent believes that it is an attractive position then at the next time point after t the agent will go to position p'

Here, b3(p,4) represents the belief that food was twice present at p, and subsequently disappeared (in other words, a belief that p is an attractive position, since food might show up again). Note that, although the mechanism described here is quite different from, e.g., machine learning, this type of behaviour nevertheless can be qualified as learning behaviour. The reason for this is that the behaviour can be split into two distinct phases: one in which nothing was learned, and one in which the agent has learned which positions are useful by maintaining a history of previous observations.

Logical relations
ES4 ≡ EP1 & EP2 & EP3 & EP4 & EP5 & ((EP6 & EP7 & EP10 & EP11 & EP12)
OR (EP6 & EP7 & EP8 & EP11 & EP12 & EP13) OR (EP9 & EP14))
BS4 ≡ BP2 & BP3 & BP4 & BP5
CS4 ≡ BP2 & CP1 & CP2 & CP3 & CP4 & CP5 & CP6 & CP7 & CP8 & CP9 & CP10 &
CP11 & CP12 & CP13 & CP14 & CP15
BS4 & ES4 ⇒ VP
CS4 ⇒ BS4

6 Formalisation of Temporal Complexity

The Complexity Monotonicity Thesis discussed earlier involves environmental, behavioural and cognitive dynamics of living systems. In Section 2 it was shown that based on a given complexity measure cm this thesis can be formalised by:

$$\mathrm{cm}(E_1) \leq \mathrm{cm}(E_2) \;\Rightarrow\; \mathrm{cm}(B_1) \leq \mathrm{cm}(B_2) \;\;\&$$
$$\mathrm{cm}(B_1) \leq \mathrm{cm}(B_2) \;\Rightarrow\; \mathrm{cm}(C_1) \leq \mathrm{cm}(C_2)$$

What remains is the existence or choice of the complexity measure function cm. To measure degrees of complexity for the three aspects considered, a temporal perspective is chosen: complexity in terms of the temporal relationships describing them. For example, if references have to be made to a larger number of events that happened at different time points in the past, the temporal complexity is higher. The temporal relationships have been formalised in the temporal language TTL based on predicate logic. This translates the question how to measure complexity to the question how to define complexity of syntactical expressions in such a language. In the literature an approach is available to define complexity of expressions in predicate logic in general by defining a function that assigns a *size* to every expression [7]. To measure complexity, this approach was adopted and specialised to the case of the temporal language TTL. Roughly spoken, the complexity (or size) of an expression is (recursively) calculated as the sum of the complexities of its components plus 1 for the composing operator. In more details it runs as follows.

Similarly to the standard predicate logic, predicates in the TTL are defined as relations on terms. The size of a TTL-term t is a positive natural number s(t) recursively defined as follows:

(1) s(x)=1, for all variables x.

(2) s(c)=1, for all constant symbols c.

(3) $s(f(t1,\dots, t_n))= s(t1) + \dots + s(t_n) + 1$, for all function symbols f.

For example, the size of the term observed(not(at(food(x), p))) from the property BP1 (see the Appendix) is equal to 6.

Furthermore, the size of a TTL-formula ψ is a positive natural number $s(\psi)$ recursively defined as follows:

(1) $s(p(t_1,\dots, t_n))= s(t_1) + \dots + s(t_n) +1$, for all predicate symbols p.

(2) $s(\neg\varphi)=s((\forall x)\ \varphi)= s((\exists x)\ \varphi) = s(\varphi)+1$, for all TTL-formulae φ and variables x.

(3) $s(\varphi\&\chi) = s(\varphi|\chi) = s(\varphi\Rightarrow\chi) = s(\varphi)+ s(\chi)+1$, for all TTL-formulae φ, χ.

In this way, for example, the complexity of behavioural property BP1 amounts to 53, and the complexity of behavioural property BP2 is 32. As a result, the complexity of the complete behavioural specification for the stimulus-response case (which is determined by BP1 & BP2) is 85.

Using this formalisation of a complexity measure as the size function defined above, the complexity measures for environmental, internal cognitive, and behavioural dynamics for the considered cases of stimulus-response, delayed response, goal-directed and learning behaviours have been determined. Table 1 provides the results (see the Appendix for all properties).

Table 1. Temporal complexity of environmental, behavioural and cognitive dynamics.

Case	Environmental dynamics	Behavioural dynamics	Cognitive dynamics
Stimulus-response	262	85	85
Delayed response	345	119	152
Goal-directed	387	234	352
Learning	661	476	562

The data given in Table 1 confirm the Complexity Monotonicity Thesis put forward in this paper, that the more complex the environmental dynamics, the more complex the types of behaviour an agent needs to deal with the environmental complexity, and the more complex the behaviour, the more complex the internal cognitive dynamics.

7 Discussion

In this paper, the temporal complexity of environmental, behavioural, and cognitive dynamics, and their mutual dependencies, were explored. As a refinement of Godfrey-Smith's Environmental Complexity Thesis [4], the Complexity Monotonicity Thesis was formulated: for more complex environments, more complex behaviours are needed, and more complex behaviours need more complex internal cognitive dynamics. A number of example scenarios were formalised in a temporal language, and the complexity of these formalisations was measured. Complexity of environment, behaviour and cognition was taken as temporal complexity of dynamics of these three aspects, and the formalisation of the measurement of this temporal complexity was based on the complexity of the syntactic expressions to characterise these dynamics in a predicate logic language, as known from, e.g., [7]. The outcome of this approach is that the results support the Complexity Monotonicity Thesis.

Obviously, the results as reported in this paper are no generic proof for the correctness of the Complexity Monotonicity Thesis. Instead, the paper should rather be seen as a case study in which the thesis was tested positively. However, the approach taken for this test was not completely arbitrary: the used complexity measure is one of the standard approaches to measure complexity of syntactical expressions [7]. Moreover, the formal specifications were constructed very carefully, to ensure that no shorter specifications exist that are equivalent. Although no formal proof is given that the used specifications are indeed the shortest possible ones, the construction of these specifications has been an iterative process in which multiple authors have participated. To represent the specifications, the language TTL was just used as a vehicle. Various similar temporal languages could have been used instead, but we predict that this would not significantly influence the results.

Nevertheless, there are a number of alternative possibilities for measuring complexity that might in fact influence the results. Among these is the option to use complexity measures from information theory based on the amount of entropy of a system, such as [1]. In future work, such alternatives will be considered as well. Another challenging direction for future work is the possibility to establish a uniform approach for specification of dynamic properties for environment, behaviour, and cognition. Such an approach may, for example, prescribe a limited number of predefined concepts that can be used within the dynamic properties.

Another issue that is worth some discussion is the fact that the Complexity Monotonicity Thesis can also be considered in isolation of Godfrey-Smith's Environmental Complexity Thesis. Although it was used as a source of inspiration to explore for the more refined Complexity Monotonicity Thesis, the Environmental Complexity Thesis as such was not investigated in this paper. Doing this, again from an agent-based modelling perspective, is another direction for future work. To this end, techniques from the area of Artificial Life may be exploited, e.g., to perform social simulations and observe whether more complex agents evolve in a way that supports the Environmental Complexity Thesis.

In [4], in particular in Chapters 7 and 8, mathematical models are discussed to support the Environmental Complexity Thesis, following, among others [9] and [12]. These models are made at an abstract level, abstracting from the temporal dimension of the behaviour and the underlying cognitive architectures and processes. Therefore, the more detailed temporal complexity as addressed in this paper is not covered. Based on the model considered, Godfrey-Smith [4] concludes that the flexibility to accommodate behaviour to environmental conditions, as offered by cognition, is favoured when the environment shows (i) unpredictability in distal conditions of importance to the agent, and (ii) predictability in the links between (observable) proximal and distal. This conclusion has been confirmed to a large extent by the formal analysis described in this paper. Comparable claims on the evolutionary development of learning capabilities in animals are made in work such as [13] and [10]. According to these authors, learning is an adaptation to environmental change. All these are conclusions at a global level, compared to the more detailed types of temporal complexity considered in our paper, where cognitive processes and behaviour extend over time, and their complexity can be measured in a more detailed manner as temporal complexity of their dynamics.

References

1. Berlinger, E. (1980). An information theory based complexity measure. In Proceedings of the Natural Computer Conference, pp. 773-779.
2. Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J. (2006). Specification and Verification of Dynamics in Cognitive Agent Models. In: *Proceedings of the Sixth International Conference on Intelligent Agent Technology, IAT'06*. IEEE Computer Society Press, 2006, pp. 247-254.
3. Darwin, C. (1871). *The Descent of Man*. John Murray, London.
4. Godfrey-Smith, P., (1996). *Complexity and the Function of Mind in Nature*. Cambridge University Press.
5. Hills, T.T. (2006). Animal Foraging and the Evolution of Goal-Directed Cognition. *Cognitive Science*, vol. 30, pp. 3-41.
6. Hunter, W.S. (1912). The delayed reaction in animals. *Behavioral Monographs*, 2, 1912, pp. 1-85
7. Huth, M. and Ryan, M. (2000). Logic in Computer Science: Modelling and reasoning about computer systems, Cambridge University Press.
8. McMillan, K.L. (1993). Symbolic Model Checking: An Approach to the State Explosion Problem. PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, 1992. Published by Kluwer Academic Publishers, 1993.
9. Moran, N. (1992). The evolutionary maintenance of alternative phenotypes. *American Naturalist*, vol. 139, pp. 971-989.
10. Plotkin, H. C. and Odling-Smee, F. J. (1979). Learning, Change and Evolution. *Advances in the Study of Behaviour 10*, pp. 1-41.
11. Sharpanskykh, A., Treur, J. (2006). Verifying Interlevel Relations within Multi-Agent Systems. In: *Proceedings of the 17th European Conference on Artificial Intelligence, ECAI'06*. IOS Press, 2006, pp. 290-294.
12. Sober, E. (1994). The adaptive advantage of learning versus a priori prejustice. In: *From a Biological Point of View*. Cambridge University Press, Cambridge.
13. Stephens, D. (1991). Change, regularity and value in evolution of animal learning. *Behavioral Ecology*, vol. 2, pp. 77-89.
14. Tinklepaugh, O.L. (1932). Multiple delayed reaction with chimpanzees and monkeys. *Journal of Comparative Psychology*, 13, 1932, pp. 207-243.
15. Vauclair, J. (1996). *Animal Cognition*. Harvard Univerity Press, Cambridge, MA.
16. Wilson, O. (1992). *The Diversity of Life*. Harvard University Press, Cambridge, MA.