A Computational Model of Habit Learning to Enable Ambient Support for Lifestyle Change

Michel C.A. Klein, Nataliya Mogles, Jan Treur, Arlette van Wissen

VU University Amsterdam, Agent Systems Research Group
De Boelelaan 1081, 1081HV Amsterdam, The netherlands
`{michel.klein|nm.mogles|treur|wissen}@few.vu.nl`

Abstract. Agent-based applications have the potential to assist humans in their lifestyle change, for instance eliminating addictive behaviours or adopting new healthy behaviours. In order to provide adequate support, agents should take into consideration the main mechanisms underlying behaviour formation and change. Within this process habits play a crucial role: automatic behaviours that are developed unconsciously and may persist without the presence of any goals. Inspired by elements from neurological literature, a computational model of habit formation and change was developed as a basis for support agents able to assist humans in lifestyle and behaviour change. Simulations are presented showing that the model exhibits realistic human-like behaviour.

Keywords: habit learning, computational agent model, lifestyle change support

1 Ambient Support for Lifestyle Change

In Western societies health policy is directed at the reduction of medical costs by switching more and more from the treatment of diseases resulting from unhealthy lifestyle to promotion of healthy lifestyle habits [18]. Lifestyle change may comprise eliminating bad habits, for example addictive behaviours (e.g., smoking, alcohol or drugs use), and may simultaneously be directed at adopting new healthy habits, such as dieting and increasing physical activity (e.g., [13], [20], [24], [29], [30]). Considering the fact that lifestyle change requires intensive support, monitoring and supervision (e.g., [25]), the potential of smart ambient applications that assist humans in their daily life is substantial, as they allow for constant monitoring and instant feedback. In order to provide adequate support for humans, these support applications should be able to reason about the main determinants of human behaviour and the mechanisms underlying behaviour change.

Apart from conscious goals and decisions, human behaviour is often based on *habits* – automatic behaviours that can be developed and maintained unconsciously. Habits may persist without the presence of any clear and definite goals and are very difficult to overcome. The model for habit learning and change presented In this paper can be used as basis for ambient intelligence applications to support lifestyle change. Using the model, an ambient application can predict – given a certain context or cue – the behaviour of a person with already formed habits. In addition, the system can

exploit the model to reason about required changes in the context or goals that need more attention in order to form new habits or get rid of old ones.

The proposed computational model of habit learning was inspired by elements from the neurological literature on habit learning (e.g., [1], [8], [10], [31]), and neural plasticity, such as Hebbian learning (e.g., [3], [11]), and adopts such adaptive mechanisms. The model has been formally specified in an executable manner, in order to conduct experiments and to allow the model to be embedded in an intelligent software agent that can support humans in their lifestyle and behaviour change.

This paper is organized as follows. Section 2 addresses some background information on habit learning and change and the neural mechanisms underlying these processes. In Section 3 the description of the model is presented, Section 4 demonstrates some simulation results. Automated verification of the model is presented in Section 5, and finally, Section 6 contains discussion on the topic.

2   Background on Habit Learning

Habits are learned dispositions to repeat past responses (cf. [31], p. 843), which by themselves were goal-driven. Once habits have been acquired, they are triggered by the cues in a certain context that co-occurred frequently with past performance, and which activate habitual responses directly, without the mediation of goals. These cues can be locations, objects, sequence of actions or presence of particular persons during or preceding the action performance. Habits formation corresponds thus to a context-response learning that is acquired slowly with experience ([31], p. 844).

Behaviourists described habits as behaviour as creation of connections between stimulus and a particular response (e.g., [26], [28]). The cognitivist perspective on human behaviour suggests the existence of a central executive controller of behaviour (e.g., [21]). Nowadays, neurological literature describes the mechanisms underlying habit formation, which explain the behaviourists' stimulus-response-based learning phenomenon, and introduces the concept of neural plasticity. Learning occurs due to the change of the connections strengths, for example, based on a principle known as Hebbian learning (e.g., [3], [11], [16]). It states that if two or more neurons are co-activated, the connections between these neurons strengthen. For example, repeated action in a certain context results to the gradual strengthening of the connection between the context representation and this particular response.

These associations are difficult to override, though it is possible to influence habits (indirectly) via the activation of new goals. Strong goals that aim to direct one's behaviour are associated with activation in the prefrontal cortex. This activation can inhibit the activation of subcortical structures (e.g., basal ganglia and cerebellum), associated with habitual behavior (e.g., [1], [8], [19], [27], [28]). Thus, when habits and goals are both present to guide action, they interact such that under some circumstances humans respond habitually and under other they exert regulatory control to inhibit the cued response.

Although a habit is no longer goal-mediated, it can be regulated by post-hoc goal inference or cue control, for example, by 1) inhibiting the performance of responses, 2) drawing one's attention to the undesired behavior, 3) associating the learned context with multiple responses or 4) altering exposure to the cues in the context [31].

Summarising, from neurological literature such as [1], [8], [10], [15], [19], [23], [27], [32], [36] the following characteristics of habit learning have been identified:

1. Under repeated occurrence of cues and under influence of goal-directed behaviour leading to satisfaction, habits are developed.
2. When a habit has developed, the behaviour will also occur without the presence of a goal, when the cue is present.
3. A developed habit will persist when the relevant goal is present, also in absence of the cue.
4. When a habit was developed based on a goal, and this goal is changed to another (competitive) goal, then the habit can change to a new habit.

These patterns have served as requirements for the design of the adaptive computational model described in Section 3. The patterns themselves will be formalized in Section 4 and checked for simulation traces of the computational model.

3 The Computational Model for Habit Learning

The structure of the computational model presented in this section is based on the literature described in the previous section. The model is at a cognitive level, which still reflects the underlying neurological concepts, but without taking into account too many neurological details. It uses temporal relationships to describe the mechanisms at work. An overview of the model is depicted in Fig. 1. It enables two alternative ways (paths) in which behaviour can be generated. The first is by the activation of a long term goal (e.g., loose weight), a short term goal corresponding to this long term goal (reduce consumption of high calorie food), generation of an intention (able to achieve the goals), and finally execution of this intended action. The second path goes directly via cue activation in a certain context to the activation of a particular intention that leads to the action execution. This path corresponds to the habit, which is learned over time: the connection between cue and intention changes dynamically after their simultaneous activation according to the Hebbian learning principle.

In the model also the influence of feeling on the chosen action has been incorporated: frequent execution of a particular action provides a reinforcement by the feeling of satisfaction after the performed action, and this feeling leads in turn to the higher activation of the intention related to this action. For example, a positive feeling of satisfaction resulting from the consumption of delicious cookies will lead to higher activation of the intention of eating these cookies. The model allows for multiple goals and intentions that result in behaviour. In principle each long term goal has connections with different strengths to short term goals, and the same holds for cues.

The dynamical relationships below describe the model in semi-formal form and in a formal temporal relation notation in LEADSTO (cf. [6]). Within LEADSTO a dynamic property or temporal relation $a \twoheadrightarrow b$ denotes that when a state property a (or conjunction thereof) occurs, then after a certain time delay, state property b will occur. Below, this delay will be taken as a uniform time step Δt. The first dynamic relationship addresses the Hebbian learning principle applied for the connections between cues en intentions, as also described in ([11], p. 406).

LP1 Cue-intention connection adaptation
If   relevant cue C with level V_1 occurs and intention I has value V_2
and   learning rate from cue C to intention I is η and   extinction rate from cue C to intention I is ζ

and the connection strength between cue C and intention I is w_1

then after Δt the connection from cue C to intention I will have
strength $w_1 + (\eta * V_1 * V_2(\ 1 - w_1) - \zeta * w_1) * \Delta t$

cue(C, V1) & intention(I, V2) & learning_rate(C, I, η) & extinction_rate(C, I, ζ) & connection_strength(C, I, w1)

→ connection_strength(C, I, w1 + (η*V1*V2(1 - w1) - ζ *w1)* Δt)
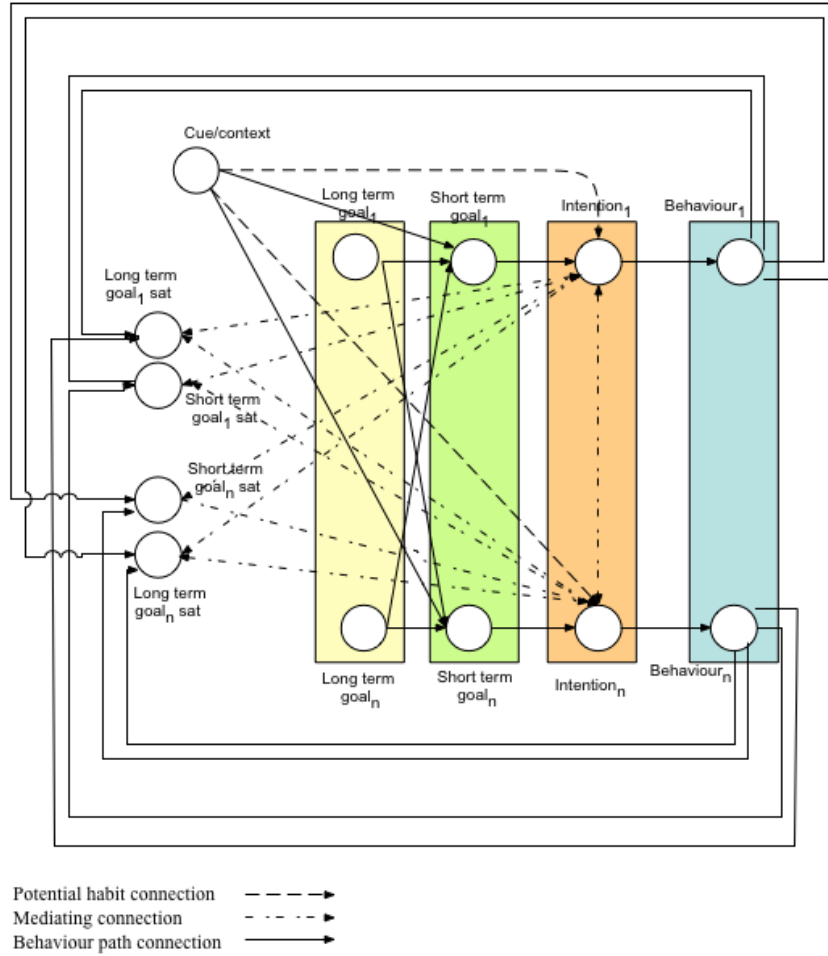
Fig. 1. Computational model for habit learning: overview

The following relationship specifies how activations of short term goals are determined based on long term goals and cues.

LP2 Short term goal from cue and long term goals

If relevant cue C with level V_0 occurs,

and long term goal LG_1 has value V_1 … and long term goal LG_n has value V_n

and the connection strength between cue C and short term goal SG is w_0

and the connection strength between long term goal LG_1 and short term goal SG is w_1

and the connection strength between long term goal LG_2 and short term goal SG is w_2

...
and the connection strength between long term goal LG_n and short term goal SG is w_n
and short term goal SG_1 has value V_3
then short term goal SG_1 after Δt will have
level $V_3 + \alpha(g(\sigma_1,\tau_1, V_0, V_1, V_{2,\dots,} V_n, w_0, w_1, w_2, \dots, w_n) - V_3)\,\Delta t$

cue(C, V0) & ltgoal(LG1, V1) & … ltgoal(LGn, Vn) & connection_strength(C, SG, w0) &
connection_strength(LG1, SG, w1) & … connection_strength(LGn, SG, wn) & stgoal(SG1, V3)

→ stgoal(SG1, V3 + α(g(σ₁, τ₁, V0, V1 , V2, Vn, w0, w1, w2, wn) − V3) Δt)

Here α is a speed parameter that defines the impact of long term goals and context cues upon the new activation value of the short term goal. Moreover, g is a combination function for which various choices are possible; a logistic threshold function has been chosen: $g(\sigma, \tau, V_0, \dots, V_n, w_0, \dots, w_n) = th(\sigma, \tau, w_0V_0 + \dots + w_nV_n)$ with

$th(\sigma, \tau, V) = 1/(1 + e^{-4\sigma(V-\tau)})$

Parameters σ and τ define steepness and threshold values of the function. The threshold function ensures that the value of the goal is most often either close to zero or close to one. Only when the input for the threshold function is close to the threshold value itself, the values of the goal are somewhere between *0* and *1*. In all subsequent formulae the combination function g is always based on a threshold function of this form.

The third relationship of the model describes how intentions are determined. Intentions depend on short term goals and cues, and the feelings of satisfaction for both short and long term goals. Moreover, different intentions also affect each other by a form of mutual inhibition. Note that for the sake of simplicity in LP3 only two long term and short term goals are considered.

LP3 Intention dynamics
If short term goal satisfaction $SGSAT_1$ has value V_1
and long term goal satisfaction $LGSAT_1$ has value V_2
and short term goal satisfaction $SGSAT_2$ has value V_3
and long term goal satisfaction $LGSAT_2$ has value V_4
and relevant short term goal SG_1 has value V_5
and relevant cue C has value V_6
and intention I_1 that corresponds to these goals has value V_7
and intention I_n has value V_8
and the connection strength between intention I_1 and intention I_n is w_1
and the connection strength between short term goal SG_1 and intention I_1 is w_3
and the connection strength between cue C and intention I_1 is w_4
and the connection strength between short term goal satisfaction $SGSAT_1$ and intention I_1 is w_5
and the connection strength between long term goal satisfaction $LGSAT_1$ and intention I_1 is w_6
and the connection strength between short term goal satisfaction $SGSAT_2$ and intention I_1 is w_7
and the connection strength between long term goal satisfaction $LGSAT_2$ and intention I_1 is w_8
then intention I_1 that corresponds to these goals after Δt will have value
$V_7 + \beta (g((\sigma_2,\tau_2, V_8, V_9, V_1, V_2, V_3, V_4, V_5, V_6, w_3, w_4, w_5, w_6, w_7, w_8, w_1) - V_7) * \Delta t$

stg_satisfaction(SGSAT1, V1) & ltg_satisfaction(LGSAT1, V2) & stg_satisfaction(SGSAT2, V3) &
ltg_satisfaction(LGSAT2, V4) & stgoal(SG1, V5) & cue(C, V6) & intention(I1, V7) & intention(In, V8) &
connection_strength(I1, In, w1) & connection_strength(SG1, I1, w3) & connection_strength(C, I1, w4) &
connection_strength(SGSAT1, I1, w5) & connection_strength(LGSAT1, I1, w6) &
connection_strength(SGSAT2, I1, w7) & connection_strength(LGSAT2, I1, w8)

→ intention(I1, V7 + β (g(σ₂, τ₂, V8, V9, V1, V2 , V3, V4, V5, V6, w3, w4, w5, w6, w7, w8, w1) − V7) Δt)

Here β is a parameter that defines the impact of inhibition of other intentions, and the feeling of satisfaction from the performed actions upon the intention to perform new actions. Weight w_1 is negative here as it defines inhibition from the alternative competing intention(s). It is assumed that different intentions are conflicting, in other words one cannot perform two behaviour simultaneously to satisfy different goals; for this reason the weights between the intentions are always negative, or inhibitory. The step from intention to behaviour has been kept simple:

LP4 From intention to behaviour
If   intention I with level V occurs,   and $V >$ *threshold*
then   behaviour with level V will occur

intention(I, V) & V> threshold $\rightarrow$ behaviour(B, V)

The feeling of satisfaction for a long term goal was modelled as follows:

LP5 Long term goal satisfaction
If   behaviour B_1 with level V_1 occurs   and   intention I_1 has value V_2
and   long term goal LG corresponding to this behaviour has value V_3
and   long term goal satisfaction *LTSAT* has value V_4
and   connection strength from behaviour B_1 to the long term goal satisfaction *LTSAT* is w_1
and   connection strength from intention I_1 to long term goal satisfaction *LTSAT* is w_2
then   long term goal satisfaction *LTSAT* after Δt will be $V_4 + \theta\ (f((\sigma_3, \tau_3, V_3 , V_1, V_2,\ w_1, w_2) - V_4) * \Delta t$

behaviour(B1, V1) & intention(I1, V2) & ltgoal(LG, V3) & ltg_satisfaction(LTSAT, V4) &
connection_strength(B1, LTSAT, w1) & connection_strength(I1, LTSAT, w2)
$\rightarrow$ ltg_satisfaction(LTSAT, V4 + θ (f(σ₃, τ₃, V3 , V1, V2, w1, w2) − V4) * Δt)

Here parameter θ defines the impact of a long term goal, behaviours and intentions upon the long term goal satisfaction. The feeling of satisfaction for a short term goal was modelled in a similar manner:

LP6 Short term goal satisfaction
If   behaviour *B1* with level V_1 occurs   and   intention I_1 has value V_2
and   short term goal SG corresponding to this behaviour has value V_3
and   short term goal satisfaction *STSAT* has value V_4
and   connection strength from behaviour B_1 to the short term goal satisfaction *STSAT* is w_1
and   connection strength from intention I_1 to short term goal satisfaction *STSAT* is w_2
then   short term goal satisfaction *STSAT* after Δt will
        be $V_4 + \theta\ (f((\sigma_4, \tau_4, V_3 , V_1, V_2,\ w_1, w_2) - V_4)\ \Delta t$

behaviour(B1, V1) & intention(I1, V2) & stgoal(SG, V3) & stg_satisfaction(STSAT, V4) &
connection_strength(B1, STSAT, w1) & connection_strength(I1, STSAT, w2)
$\rightarrow$ stg_satisfaction(STSAT, V4 + θ (f(σ₄, τ₄, V3 , V1, V2, w1, w2) − V4) * Δt)

4   Simulation and Verification

The cognitive computational model described in the previous section was implemented in the Matlab environment. A number of simulations of 50 and 200 time steps have been performed. For the sake of simplicity only two initial long term goals and the corresponding behaviours were assumed. In this section four example simulation runs of 50 time steps are presented. These simulations illustrate the ability of the computational model to exhibit important patterns of habit learning and change. In Table 1 the values are shown used for learning and extinction rate, steepness and

threshold values, speed factors, and connection weights (note that weight values for interaction between two options are symmetric). In order to investigate whether the computational model indeed learns and behaves according to what is expected, some logical properties (requirements) have been identified, formalized, and verified against the simulation traces of the model (see also the characteristics informally described at the end of Section 2). In this section, first the language used to express such properties is briefly introduced, followed by the specification of the actual properties, presentation of an example trace illustrating the pattern, and the result of their verification.

Table 1. Parameter and connection weight values used

η	ζ	α	σ_1	τ_1	β	σ_2	τ_2	thr	θ	σ_3	τ_3	σ_4	τ_4
0.5	0.01	0.8	15	0.9	0.8	20	0.5	0.5	0.8	15	0.2	15	0.6

connection	weight	connection	weight	connection	weight
cue-intention1 initially	0.1	stsat1- intention2	0.2	intention1-stsat1	0.9
cue-intention2 initially	0.1	intention1-behaviour1	1	behaviour1-stsat1	0.9
cue-stgoal1	0.1	intention2-behaviour2	1	intention2-ltsat1	0.2
cue-stgoal2	0.1	intention1-intention2	-0.9	behaviour2-ltsat1	0.1
ltgoal1-stgoal1	0.9	intention2-intention1	-0.9	intention2-stsat1	0.2
ltgoal1-stgoal2	0.2	intention1-ltsat1	0.9	behaviour2-stsat1	0.1
ltgoal2-stgoal1	0.2	behaviour1-ltsat1	0.9	behaviour1- ltsat2	0.1
ltgoal2-stgoal2	0.9	intention2-ltsat1	0.1	behaviour1-stsat2	0.1
ltsat1- intention1	0.9	intention2-stsat1	0.1	behaviour2- ltsat2	0.9
ltsat1- intention2	0.2	intention1-ltsat2	0.1	behaviour2- stsat2	0.9
stsat1- intention1	0.9	intention1-stsat2	0.1		

Formal specification of desired properties of the computational model enables automatic verification of them against simulation traces. This was performed using the hybrid language TTL and its software environment [5]. TTL is built on atoms referring to *states* of the world, *time points* and *traces*, i.e. trajectories of states over time. *Dynamic properties* are temporal statements formulated with respect to traces in the following manner. Given a trace γ, the state in γ at time point t is denoted by state(γ, t). These states are related to state properties via the infix predicate $\models$, where state(γ, t) $\models$ p denotes that state property p holds in trace γ at time t. Based on these statements, dynamic properties are formulated in a sorted first-order predicate logic, using quantifiers over time and traces and the usual first-order logical connectives such as $\neg$, $\wedge$, $\vee$, $\Rightarrow$, $\forall$, $\exists$. For more details on TTL, see [5].

Each of the three subsections addresses one scenario. In the figures that demonstrate the simulation results, time is depicted on the horizontal axis and the activation values of the variables of interest are depicted on the vertical axis.[1]

---

[1] Note that gradual values for intentions, cues and goals are assumed, which represent the strength of the presence of these variables as perceived by the subject.

4.1 Habit formation

In this simulation a specific behaviour is generated by a strong long term goal related to this behaviour in the presence of a strong cue. As a result even after a decrease of the value of the goal corresponding to this behaviour after time point 24, the behaviour persists up to end of the simulation; see Fig. 2. The value of the second long term goal is kept low during the whole simulation; therefore the second type of behaviour that corresponds to this goal does not come to expression. To verify this pattern formally, it first has to be checked whether a specific behaviour results from the presence of a high-level goal and a cue:

P0: Long-term goal and cue leads to behaviour
If a cue and a high-level goal are present for a certain time duration MIN_DURATION, then at some later time the corresponding behaviour will be present.
∀γ:TRACE, t:TIME   [habit_learning_phase(γ:TRACE, t:TIME, MIN_DURATION, ACT_VALUE2)
⇒ ∃t2:TIME > t , R3:REAL   state(γ, t2) |= has_value(beh1,R3) & R3 > ACT_VALUE2]

Here (and in the other properties below) the following abbreviation is used:

habit_learning_phase(γ:TRACE, t:TIME, MIN_DURATION:INTEGER, ACT_VALUE:REAL) ≡
∀t2:TIME > t & t2 < t + MIN_DURATION [∃R1:REAL state(γ, t2) |= has_value(ltg1, R1) & R1 > ACT_VALUE2
& ∃R2:REAL state(γ, t2) |= has_value(cue1, R2) & R2 > ACT_VALUE2]

Fig. 2. Habit formation and persistence   **Fig. 3.** Habit formation and cue removal

Property P0 corresponds to characteristic pattern 1 of habits as defined at the end of Seciton 2, and has been automatically checked and proven to be true for the following values of the constants: MIN_DUR = 2, MAX_LEVEL_P1 = 0.5, ACT_VALUE1 = 0.5, ACT_VALUE2 = 0.7. All properties described in the remainder of this section have been automatically verified and found satisfied for these values. The pattern of habit formation itself (characteristic pattern 2) was specified as follows:

P1: Habit persistence
If a cue and a high-level goal have been present for some time period MIN_DURATION, the behaviour will exist in the presence of a cue even if the goal is no longer present.
∀γ:TRACE, t:TIME
[habit_learning_phase(γ:TRACE, t:TIME, MIN_DURATION, ACT_VALUE2) &
∀t2:TIME > t + MIN_DURATION, R3: REAL [state(γ, t2) |= has_value(ltg1, R3) ⇒ R3 < ACT_VALUE1]]
⇒ [∀t3:TIME > t2, R4: REAL   state(γ, t3) |= has_value(cue1, R4) & R4 > ACT_VALUE2
⇒ ∃t4:TIME>t3, R5:REAL state(γ, t4) |= has_value(beh1, R5) & R5 > ACT_VALUE2]]

When in the scenario in Fig. 3, after time point 26 the value of the cue is substantially decreased, habitual behaviour is not performed anymore from time point 31. As expected, and shown in Fig. 3, the second behaviour ('behaviour 2') does not occur. Formally, the illustrated characteristic is specified as follows.

P2: Habit and cue removal

If a habit is formed and the cue and the goal are no longer present, the behaviour will after some time cease to exist.

$\forall\gamma$:TRACE, t:TIME [[habit_learning_phase(γ, t, MIN_DURATION, ACT_VALUE2) &

$\forall$t2:TIME > t + MIN_DURATION, R1, R2: REAL [state(γ, t2) |= has_value(cue1,R1) $\Rightarrow$ R1 < ACT_VALUE1 &

state(γ, t2) |= has_value(ltg1, R2) $\Rightarrow$ R2 < ACT_VALUE1]

$\Rightarrow$ $\exists$t3:TIME > t2, R3: REAL state(γ, t3) |= has_value(beh1,R3) & R3 < ACT_VALUE1]]

### 4.2   Influence of long term goal on behaviour

This scenario demonstrates how behaviour is influenced by goals in the absence of the learned cue; see Fig. 4. In the beginning habitual behaviour is formed: a strong cue is present and a behaviour pattern that coincides with the first long term goal. The value of the goal remains the same during the whole run, but the cue almost disappears after time point 24. The low value of the cue does not prevent the behaviour to occur due to the strong influence of the long term goal. This corresponds to characteristic habit pattern 3 from Section 2. This was specified as follows.

P3: Habit and cue removal in presence of strong goal

If a habit is formed, the behaviour will still exist if the cue is not present any more and the high-level goal is present.

$\forall\gamma$:TRACE, t:TIME,

[[habit_learning_phase(γ, t, MIN_DURATION, ACT_VALUE2) &

  $\forall$t2:TIME > t + MIN_DURATION, R3: REAL   [state(γ, t2) |= has_value(cue1, R3) $\Rightarrow$ R3 < ACT_VALUE1]

$\Rightarrow$ [$\forall$t3:TIME > t2, R4: REAL

          state(γ, t3) |= has_value(ltg1, R4) & R4 > ACT_VALUE2 $\Rightarrow$

          $\exists$t4:TIME>t3,R5:REAL state(γ, t4) |= has_value(beh1, R5) & R5 > ACT_VALUE2]]

Fig. 4. Influence of the goal on behaviour in the absence of the original cue

Fig. 5. Behaviours resulting from the goal change

4.3 The effect of goal change

In this simulation the result of the switch from one goal to another is demonstrated in the presence of a strong cue. As shown in Fig. 5, the habitual behaviour ('behaviour 1' in the picture) does not disappear immediately after adopting a new goal, conflicting with the previous one. It takes a little time to perform new behaviour pattern after the new goal has been adopted. This simulation demonstrates how the old undesired habitual behaviour can be substituted with the new positive behaviour. The formal specification of this pattern is:

P4: New goal results in new habit
If a habit is formed for long term goal ltg1, which disappears, a new behaviour will be developed if another long term goal ltg2 is present

$\forall \gamma$:TRACE, t:TIME _DURATION [[habit_learning_phase(γ, t, MIN_DURATION, ACT_VALUE2) &
  $\forall$t2 > t+MIN_DURATION, R1, R2, REAL   [state(γ, t2) |= has_value(ltg1,R1) $\Rightarrow$ R1 < ACT_VALUE1] &
          [state(γ, t2) |= has_value(ltg2,R2) $\Rightarrow$ R2 > ACT_VALUE2]
$\Rightarrow \exists$t3:TIME > t2 , R3, R4, REAL [state(γ, t3) |= has_value(beh1,R3) & R3 < ACT_VALUE1 &
          [state(γ, t3) |= has_value(beh2,R4) & R4 > ACT_VALUE2]]

Fig. 5 also shows the effects of P1, which demonstrate that the new behaviour results in a habit after some amount of time: the behaviour persists even after the corresponding long term goal is no longer present. Combined, P1 and P4 account for characteristic habit pattern 4.

5   Discussion and Conclusions

The cognitive computational model presented above can form the basis of an intelligent ambient support application. To this end, an agent based approach for creating ambient intelligence applications can be used [4]. Within such a framework, the ambient system consists of components, i.e., agents, that have context awareness about human behaviours and states, and (re)acts on these accordingly. For this purpose, the behaviour of the subject of the system (a person taken care of) relevant to the support provided should be explicitly described, e.g., via a computational model. If this is the case, an ambient agent can (re)act by undertaking actions in a knowledgeable manner that improve the human's wellbeing and performance.

Reasoning using an explicit model of the behaviour of a process is called model-based reasoning [22]. Basically, there are two ways in which model-based reasoning on habits can be used within an intelligent support application. First, predictions can be made of what will happen given certain cues /contexts, long term goals and short term goals. For example, if the system has identified a specific behaviour – such as the eating of cookies at work – several times in the past, and it has knowledge about the short-term and long-term goals, it can predict whether a person in the work-context will again eat a cookie. These predictions capabilities allow a support application to take action before an undesired habit actually took place. Second, the model can be used to perform analysis of the causes of the undesired behaviour and the effect of interventions on the behaviour of a person [9]. Causes of behaviour can be determined by backward abductive reasoning. For example, if an undesired behaviour is taking place, the presented computational model can be used to find hypothetical causes for this behaviour, for example a short term goal that leads to the

intentions for the undesired behaviours. Symmetrically, forward deductive model-based reasoning derives the effect of interventions on the behaviour. For example, determining the effect of a different or more important long term goal after some time. This can be used by the ambient intelligence application to explicitly change the situation, e.g. removing cues, generating additional intention for long term goals leading to different behaviour, or suggesting actions to create new (more desired) habits.

Existing models of habit learning take either the perspective of behaviourism that does not follow the internal mechanisms underlying habit development (e.g., [7], [17]) or propose the description of habit learning in a very detailed manner at the lowest neurological level (e.g., [2], [7], [14]). The proposed computational model is at a cognitive level, between the neurological and behavioural level. The proposed way of modeling is a manner to exploit within the computational modeling area principles from the neurological literature, by lifting neurological knowledge to a mental (cognitive/affective) level. In order to successfully model more complex and human-like behaviour, for example incorporating mutual cognitive/affective interactions, and adaptive behaviour, the modeler has to consider such numerical modeling techniques; see also [23].

In future work, the model will be deployed on actual data and used to improve habit performance. Also, the model could be improved by taking into account the environment in which a person is embedded, which is currently limited to perceiving cues, but preferably also incorporates socio-environmental factors shown to play a role in habit formation development (e.g., [12]).

References

1. Ashby, F.G., Turner, B.O., Horvitz, J.C.: Cortical and basal ganglia contributions to habit learning and automaticity. Trends Cogn Science 14, 208--215 (2010)
2. Baldassarre, G.: A modular neural-network model of the basal ganglia's role in learning and selecting motor behaviours. Cogn. Systems Research 3, 5--13 (2002)
3. Bi, G., Poo, M.: Synaptic modification by correlated activity: Hebb's postulate revisited. Annu. Rev. Neurosci. 24, 139--166 (2001)
4. Bosse, T., Hoogendoorn, M., Klein, M.C.A., Treur, J.: An Agent-Based Generic Model for Human-Like Ambience. Iin: M. Mühlhäuser, A. Ferscha, and E. Aitenbichler (eds.), Constructing Ambient Intelligence: AmI-07 Workshops Proceedings. Communications in Computer and Information Science (CCIS), vol. 11, pp. 93-103 Springer Verlag (2008)
5. Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., Treur, J.: Specification and Verification of Dynamics in Agent Models. International Journal of Cooperative Information Systems 18, 167--193 (2009)
6. Bosse, T., Jonker, C.M., Meij, L. van der, Treur, J.: A Language and Environment for Analysis of Dynamics by Simulation. International Journal of Artificial Intelligence Tools 16, 435--464 (2007)
7. Cohen, M.X., Frank, M.J.: Neurocomputational models of basal ganglia function in learning, memory and choice. Behavioural Brain Research 199, 141–156 (2009)
8. De Wit, S., et al.: Differential engagement of the ventromedial prefrontal cortex by goal-directed and habitual behaviour toward food pictures in humans. J. Neurosci. 29, 11330--11338 (2009)
9. Duell, R., Hoogendoorn, M., Klein, M.C.A., Treur, J.: An Ambient Intelligent Agent Model using Controlled Model-Based Reasoning to Determine Causes and Remedies for

Monitored Problems. In: Li, Y., Pasi, G., Zhang, C., Cercone, N., Cao, L. (eds.), Proceedings of the Second International Workshop on Human Aspects in Ambient Intelligence, HAI'08, pp. 489--494. IEEE Computer Society Press (2008)

10. Everitt, B.J., Belin, D., Economidou, D., Pelloux, Y., Dalley, J.W., Robbins, T.W.: Neural mechanisms underlying the vulnerability to develop compulsive drug-seeking habits and addiction. Philos Trans R Soc Lond B Biol Sci 363, 3125--3135 (2008)

11. Gerstner, W., Kistler, W.M.: Mathematical formulations of Hebbian learning. Biol. Cybern. 87, pp. 404–415 (2002)

12. Giles-Corti, B., Donovan, R.J.: The relative influence of individual, social and physical environment determinants of physical activity. Social Science & Medicine 54, 1793--1812 (2002)

13. Gollwitzer, P.M., Sheeran, P.: Implementation intentions and goal achievement: A meta-analysis of effects and processes. Advances in Exp. Social Psychology 38, 69--120 (2006)

14. Gurney, K., Prescott, T.J., Wickens, J.R. Redgrave, P.: Computational models of the basal ganglia: from robots to membranes. Trends Neurosci. 27, 453--459 (2004)

15. Haigh, K.Z., Kiff, L.M., Myers, J., Guralnik, V., Geib, W.C., Phelps, J., Wagner, T.: The Independent LifeStyle Assistant (I.L.S.A.): AI Lessons Learned. Proc. of the 16th Conference on Innovative Applications of AI, pp.852--857. San Jose, California (2004)

16. Lymberis, A. De Rossi, D.E.: Wearable eHealth Systems for Personalised Health Management: State of the Art and Future Challenges. IOS Press (2004)

17. Machado, A.: Learning the temporal Dynamics of Behaviour. Psychological Review 104, 241--265 (1997)

18. Miller, E.K., Cohen, J.D.: An Integrative Theory of prefrontal Cortex Function.Annual Review of Neuroscience 24, 167--202 (2001)

19. Mowrer, O.H.: Learning Theory and the Symbolic Processes. New York: Wiley (1960)

20. Neal, D. T., Wood, W.: Automaticity in situ: Direct context cuing of habits in daily life. In E. Morsella, J. A. Bargh, & P. Gollwiter (Eds.), Oxford handbook of human action (pp. 442--457). New York Oxford University Press (2009)

21. Oulette, E. A. Wood, W.: Habit and Intention in everyday life: the Multiple Processes by Which Past Predicts Future Behaviour. Psychological Bulletin 124, 54--74 (1998)

22. Pearl, J., Verma, T.S.: A theory of inferred causation, in Principles of Knowledge Representation and Reasoning: Proc. Second Int. Conference (KR '91), 441--452 (1991)

23. Port, R., van Gelder, T.J.: Mind as Motion: Explorations in the Dynamics of Cognition. MIT Press (1995)

24. Quinn, J.A., Pascoe, A., Wood, W., Neal, D.T.: Can't control yourself? Monitor those bad habits. Pers Soc Psychol Bull 36, 499--511 (2010)

25. Skinner, B.F.: The Behaviour of Organisms: An Experimental Analysis. New York: Appleton-Century (1938)

26. Tang, C., Pawlak, A.P., Prokopenko, V., West, M.O.: Changes in activity of the striatum during formation of a motor habit. Eur J Neurosci 25, 1212--27 (2007)

27. Watson, J.B.: Psychology as Behaviourist Views It. Psy. Review 20, 158--177 (1913)

28. Webb, T. L., Sheeran, P., Luszczynska, A.: Planning to break unwanted habits: Habit strength moderates implementation intention effects on behaviour change. British Journal of Social Psychology 48, 507--523 (2009)

29. Webb, T. L., Sheeran, P.: Mechanisms of implementation intention effects: The role of goal intentions, self-efficacy, and accessibility of plan components. British Journal of Social Psychology 47, 373--395 (2008)

30. Wood, W., Neal, D.T.: A New Look at Habits and the Habit-Goal Interface. Psychological Review 114, 843--863 (2007)

31. Yin, H.H., Knowlton, B.J.: The role of Basal ganglia in habit formation. Nature Reviews Neuroscience 7, 464--476 (2006)