Combining rational and biological factors in virtual agent making

Tibor Bosse · Charlotte Gerritsen · Jan Treur

Abstract To enhance believability of virtual agents, this paper presents an agent-based modelling approach for decision making, which integrates rational reasoning based on means-end analysis with personal psychological and biological aspects. The agent model developed is a combination of a BDI-model and a utility-based decision model in the context of specific desires and beliefs. The approach is illustrated by addressing the behaviour of violent criminals, thereby creating a model for virtual criminals. Within a number of simulation experiments, the model has been tested in the context of a street robbery scenario. In addition, a user study has been performed, which confirms the fact that the model enhances believability of virtual agents.

T. Bosse · C. Gerritsen (✉) · J. Treur
Department of Artificial Intelligence, Vrije Universiteit
Amsterdam, De Boelelaan 1081, 1081 HV, Amsterdam, The
Netherlands
e-mail: cg@few.vu.nl

T. Bosse
e-mail: tbosse@few.vu.nl

J. Treur
e-mail: treur@few.vu.nl

1 Introduction

In recent years, human-like virtual agents are increasingly being applied in various domains [29, 32]. Typical examples are agents in (serious) games (e.g., an instructor in a naval training simulator [13], character agents in computer-generated virtual stories [10], or conversational agents (e.g., seller representatives on the internet [25]). Recently, much research has been dedicated to developing virtual agents with more realistic graphical representations. However, the behaviour of such agents is usually not very human-like. For example, although many virtual agents currently have the ability to somehow show emotions by means of different facial expressions, it is quite difficult for them to show the right emotion at the right moment. This is in conflict with the requirement of virtual agents to closely mimic human affective behavior. Several studies in Social Sciences have shown that this is an important prerequisite for an agent to increase human involvement in the virtual environment [18]. Therefore, existing systems based on IVAs are not as effective as they could be.

A known problem encountered by developers of human-like virtual agents is to create realistic *decision making* behaviour for such agents. For humans it is known that the process of decision making is determined partly by rational means-end reasoning, and partly by subjective personal biological and psychological aspects (including, for example, their motivational and emotional state, see, e.g., [12, 14, 24, 38]). On the one hand, humans have various kinds of—partly biologically determined—desires, but on the other hand, they may have to reason rationally about which desires to fulfil. For example, a person may have the desire to eat but may decide rationally not to do this because it will make him/her fat. However, if the biological desire is too strong,

the person may decide to eat nevertheless. Thus, some mechanism is used that enables humans to make decisions in situations where both rational and biological/psychological factors play a role.

This paper introduces an approach to incorporate such mechanisms within virtual agents in order to make their behaviour more human-like. In principle, the approach is generic, i.e., it can be used to model virtual agents in various types of applications (varying from serious games to virtual stories) and in various domains (varying from flight simulators to adventure games). As an illustration of the approach, in the current paper a specific domain is chosen, namely the domain of *crime*. This is an interesting case study, since this is a typical domain in which both rational decision making and biological and psychological aspects play a role. Within the area of Criminology, a longstanding debate is whether criminal behaviour is driven by a criminal's subjective, personal biological and psychological background, or is the result of a rational, calculated choice; e.g., [11, 27]. The current paper will show how the two viewpoints can be integrated, thereby creating a behavioural model for a "virtual criminal agent".

As a starting point, the criminal agent model described in [6] has been taken, which focuses specifically on violent criminals (such as psychopaths and persons with the Intermittent Explosive Disorder), and its embedding in a social context described in [5]. This model addresses action generation based on beliefs, desires and intentions (BDI), and generation of desires and beliefs in opportunities. However, for the sake of simplicity, only one action per desire was assumed in that model, so no decision making was covered involving a choice between different options for actions to fulfil a desire. The current paper extends that model with a mechanism for utility-based multi-criteria decision making (e.g., [21, 33]) within a BDI-setting. This decision model provides a formalisation of the Rational Choice Theory within Criminology; e.g. [11]. This theory as informally discussed within Criminology describes crime as an event that occurs, for example, when an offender decides to take risk breaking the law, after considering his or her own need for money, personal values or learning experiences and how well a target is protected. The criminal assesses the chances of getting caught, the expected penalty, the value to be gained by committing the act, and his or her immediate need for that value.

In the decision model introduced in this paper, this process is modelled by introducing utilities for different possible intended actions. The utility of a certain (option for an) action is then assessed according to the extent to which it fulfils the agent's desire. In this way utilities are assessed with respect to a subjective measure focusing on a specific desire, which may be affected by the subject's specific biological and psychological background. In other words, for the individual agent, rational choice means the choice to fulfil its own desires in the best possible way. Thus, the model for desire generation based on the biological and psychological factors is integrated with a rational decision model for the choice of (intended) actions.

In this paper, Sect. 2 discusses the dynamic modelling approach that was used at a global level. In Sect. 3, a brief summary from the literature is presented on the role played by biological and psychological factors in criminal behaviour. Next, the two main components of the simulation model are presented: Sect. 4 briefly presents the model to determine desires (inspired by [6]) and Sect. 5 presents the utility-based decision making model. In Sect. 6, some simulation results are shown, addressing an example street robbery scenario. Section 7 describes the results of an experiment to evaluate how the generated simulation traces are appreciated by humans. Finally, the approach and its possible applications are discussed in Sect. 8.

2 Modelling approach

Modelling the various aspects related to criminal decision making of virtual agents in an integrated manner poses some challenges. On the one hand, qualitative aspects have to be addressed, such as beliefs, desires, and intentions, certain brain deviations, and some aspects of the environment such as the presence of certain agents. On the other hand, quantitative aspects have to be addressed, such as testosterone and serotonin levels, and utilities. Moreover, the aspects have to be integrated in such a way that the resulting model can easily be implemented within a virtual agent. For example, it should be possible to connect them to concrete events and states that occur in a virtual world, such as encounters between agents, and goals of individual agents.

It was not easy to find a modelling approach that fulfilled all of the above desiderata. Within the cognitive modelling area, a number of generic architectures exist for computational modelling of cognitive processes such as attention, memory, and decision making, for example ACT-R [1], Soar [23], and Clarion [40]. These approaches have as advantages that they distinguish in some way or another—between implicit (or subconscious) and explicit (conscious, rational) processes. For this reason, the approach put forward in this article was inspired by these approaches. However, they could not be used directly, for two reasons. First, these approaches were originally designed for the purpose of representing cognitive processes in much detail (including, for instance, specific timing parameters in terms of milliseconds), and were therefore not directly applicable to virtual agents, which need modelling constructs at a much higher level of abstraction. Second, they do not provide any mechanism to model biological factors, such as testosterone and serotonin levels.

To solve the first problem, it was decided to also take some inspiration from the area of Artificial Intelligence (and, more specifically, agent-based modelling). In this area, it is quite common to model an agent's behaviour in terms of beliefs, desires, and intentions [17, 31, 35, 36]. Thus, it was decided to take the standard BDI-model [17, 36] as a basis, and extend this with a mechanism to represent implicit, subconscious processes as done in cognitive modelling architectures. However, since the second problem (the lack of constructs for biological factors) could not be solved directly by the BDI-model, the authors decided to develop a new model, thereby integrating the BDI-model with constructs for both psychological (cognitive) as well as biological factors.

To develop this new integrated model, the modelling language LEADSTO [7] turned out to be a suitable candidate. This language integrates qualitative, logical aspects and quantitative, numerical aspects. This integration allows the modeller to exploit both logical and numerical methods for analysis and simulation. The basic building blocks of LEADSTO are so-called *executable dynamic properties*, by which direct temporal dependencies between two state properties in successive states are modelled. Their format is defined as follows. Let α and β be state properties of the form 'conjunction of ground atoms or negations of ground atoms'. In LEADSTO, the notation $\alpha \twoheadrightarrow_{e,f,g,h} \beta$, 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.

Here, atomic state properties can have a qualitative, logical format, such as an expression desire(d), expressing that desire d occurs, or a quantitative, numerical format such as an expression has_value(x, v) which expresses that variable x has value v. For more details of the language LEADSTO, see [7]. As stated above, the overall simulation model has been built by composing two models (see Fig. 1 for an overview):
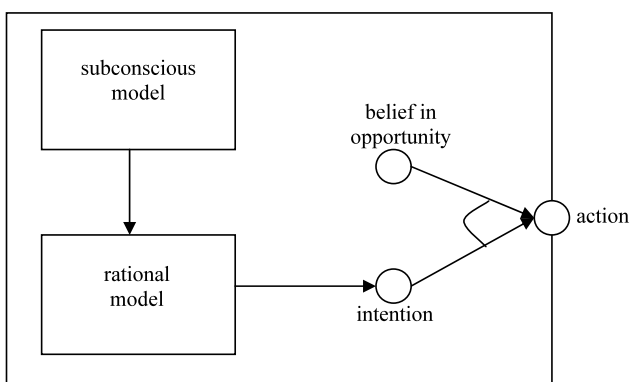

Fig. 1 Overview of criminal decision making model

1. a model to *determine desires* incorporating various biological and psychological aspects and their interactions (subconscious model)
2. a model for reasoning about *beliefs, desires and intentions*, using a BDI-model based on *utility-based decision making* (conscious, rational model).

These models have both been implemented in LEADSTO. They are described in more detail in Sects. 4 and 5, respectively.

To create the subconscious model, we first performed an extensive research in the existing literature on criminal behaviour, in collaboration with a group of experts from different disciplines. These experts, among which psychologists, criminologists and experts in AI, together agreed upon a short list of factors that were selected to be used in the model. The biological factors that were selected are based on work by Moir and Jessel [27] and by Adrian Raine [34], a well known neurocriminologist. Some of the psychological factors are inspired by work by psychologist Martine Delfos [13]. A complete overview of the factors that were selected for the model is provided in the next section.

3 Biological and psychological factors

Since the BDI model [17, 36] does not prescribe a standard way to determine how desires are created, for a particular application usually domain-specific knowledge is used. For criminal behaviour, a number of specific biological and psychological aspects seem to play a role in the generation of desires. An extensive search has been performed into literature from areas such as Criminology and Psychology (e.g., [13, 27, 34]) for aspects to be incorporated in the model. The aspects that have been selected are explained below in more detail.

A *theory of mind* of a person (e.g., [2]) describes other persons' minds by separate mental concepts, such as the person's own beliefs, desires, and intentions, and how those concepts play a role in the person's behaviour. Criminal actions are often performed by persons whose theory of mind is less developed. In recent years, more evidence is found that there often are biological reasons for this. For example, it has been found that many psychopaths have a damaged connection between the frontal lobes (concerned with conscience and remorse) and the limbic area, which generates feelings; cf. [27].

Another important aspect in crimes is *aggressiveness*. Research has pointed out that there is a correlation between aggressive behaviour and the level of testosterone. In fact, 89% to 95% of all crime is performed by males [27]. In addition, the use of alcohol or drugs may increase the violence of behaviour.

A third aspect involved in criminal behaviour is the *desire to act*, which can be related to a high level of adrenalin. If a person's adrenalin level becomes too high, (s)he somehow has to cope with this; acting decreases the adrenalin level. Thus, if the desire to act is high, then a criminal act more easily occurs. The specific types of actions that are chosen depend on another factor, the *desire to act safely*. This factor correlates with a high level of oxytocin, a hormone mainly produced by women. Persons with a high level of oxytocin have a higher tendency to cope with their desire to act by performing 'safe' actions (e.g., taking care of the 'nest') than persons with less oxytocin; they will rather perform 'less safe' actions (e.g., fighting) [13].

In addition, crimes are often committed by persons who are looking for a thrill. These persons in general have a high excitement threshold, which means that it is very difficult for them to become excited [27, 34]. As a result, they are often bored, so that they generate a *desire for actions with strong stimuli*. Such actions may become criminal actions, such as stealing, joyriding, or assaulting other people. Only by performing these actions, their desire for strong stimuli is fulfilled, and they become less bored.

Furthermore, a significant amount of committed crimes can be described as *impulsive*. They are not planned, but rather triggered by occasional opportunities. An important factor causing impulsive behaviour is a low level of blood sugar, which in turn is caused by a high insulin level and a low serotonin level [27].

A next factor that may affect the types of (criminal) actions that persons may perform, is the extent to which they have (positive or negative) feelings with respect to another person's wellbeing. When someone has a low amount of *positive feelings towards others*, (s)he does not really care about the other. Likewise, when someone has many *negative feelings towards others*, (s)he may wants to cause harm towards someone else. For example, in psychopaths, both attitudes are low: these persons hardly show any emotion concerning other persons, so for them, both the positive and the negative emotional attitude towards others are low [27].

The last two factors chosen to incorporate in the model are the *desire for high gain* and the *desire for low loss*. These concepts were chosen on the basis of the Rational Choice Theory [11]. According to this theory, to determine their actions, persons will try to minimise their expected loss or penalty (e.g., being caught, getting hurt) and maximise their gain (e.g., money, status). The theory states that criminals will make a serious decision before committing a crime, weighing pros against cons.

4 Determining desires

To determine desires, a rather complex submodel is used, incorporating dynamical system elements for the biologi-

cal and psychological aspects as mentioned earlier, varying from qualitative aspects, such as anatomical aspects concerning brain deviations (e.g., the absence of certain connections) to quantitative aspects, such as biochemical aspects concerning testosterone levels. Some example LEADSTO specifications (called Local Properties, LPs) are given below (both in informal and in formal notation):[1]

LP9 A certain level of current testosterone will lead to a corresponding level of aggressiveness.
$\forall$x:SCALE chemical_state(testosterone,current,x)
$\rightarrow\!\!\!\rightarrow_{0,0,1,1}$ desire_for_aggressiveness(x)

LP20 Observation of a stimulus with an excitement level that is lower than the excitement threshold will lead to boredom.
$\forall$s1,s2,y:INTEGER observes_stimulus(s1,s2) $\land$
excitement_threshold(y) $\land$ s2$<$y $\rightarrow\!\!\!\rightarrow_{0,0,1,1}$ boredom

LP21 Boredom leads to a high desire for actions with strong stimuli.
boredom $\rightarrow\!\!\!\rightarrow_{0,0,1,1}$ desire_for_actions_with_strong_stimuli(high)

LP29a A low blood sugar level leads to high impulsiveness.
chemical_state(blood_sugar, low)
$\rightarrow\!\!\!\rightarrow_{0,0,1,1}$ desire_for_impulsiveness(high)

The variety of biological and psychological aspects that were found relevant in the literature (such as [3, 13, 27, 34]) and are taken into account in this model, are those described in the second section above. Different combinations of these elements lead to different types of (composed) desires; e.g., the desire to perform an exciting planned nonaggressive nonrisky action that harms somebody else (e.g., a pick pocket action in a large crowd). The following LEADSTO rule generates a composed desire out of the different ingredients mentioned earlier:

LP30 A combination of values for theory of mind, desire for aggressiveness, desire to act, desire to act safely, desire for actions with strong stimuli, desire for impulsiveness, positive and negative emotional attitude towards others, and desire for high gain and low loss leads to a specific composed desire, represented as
d(has_value(theory_of_mind, s1), . . . ,
has_value(desire_for_low_loss, s10)).
$\forall$s1,s2,s3,s4,s5,s6,s7,s8,s9,s10:SCALE
theory_of_mind(s1) $\land$ desire_for_aggressiveness(s2) $\land$
desire_to_act(s3) $\land$ desire_to_act_safely(s4) $\land$
desire_for_actions_with_strong_stimuli(s5) $\land$
desire_for_impulsiveness(s6) $\land$
emotional_attitude_towards_others(pos,s7) $\land$
emotional_attitude_towards_others(neg,s8) $\land$
desire_for_high_gain(s9) $\land$ desire_for_low_loss(s10) $\rightarrow\!\!\!\rightarrow_{0,0,1,1}$
desire(d(has_value(theory_of_mind, s1), . . . ,
has_value(desire_for_low_loss, s10)))

---

[1]See [6] for the complete criminal agent model, and [5] for its embedding in a social context.

5 Utility-based reasoning about intentions

As in [6], part of the model for criminal behaviour is based on the BDI-model, which bases the preparation and performing of actions on beliefs, desires and intentions (e.g., [17, 36]). In this model an action is performed when the subject has the intention to do this action and it has the belief that the opportunity to do the action is there. Beliefs are created on the basis of stimuli that are observed. The intention to do a specific type of action is created if there is a certain desire, and there is the belief that in the given world state, performing this action will fulfill this desire. The BDI-model was specified by:

LP31 Desire d combined with the belief that a certain action a will lead to the fulfillment of that desire will lead to the intention to perform that action.
$\forall$d:DESIRE $\forall$a:ACTION desire(d) $\wedge$ belief(satisfies(a, d)) $\twoheadrightarrow_{0,0,1,1}$ intention(a)

LP32 The belief that there is an opportunity to perform a certain action combined with the intention to perform that action will lead to the performance of that action.
$\forall$a:ACTION belief(opportunity_for(a)) $\wedge$ intention(a)
$\twoheadrightarrow_{0,0,1,1}$ performed(a)

However, to assess and compare different options, and select a best option, as an extension to this basic BDI-model utilities are to be assigned and combined, addressing the degree to which an action satisfies a desire. The notion of utility to be used requires some reflection. Sometimes this may be considered a rational notion with an absolute, intersubjective (or objective) status. For two agents with a kind of standard internal functioning, considered rational, this intersubjectivity may be a reasonable assumption. However, if the internal processes are different it is less reasonable. One agent may have preferences different from those of the other agent, and hence be satisfied with a situation that is not satisfactory for the other agent. As an example, multi-attribute negotiation aims at exploiting such differences in preferences between agents in order to the benefit of both; e.g., [7, 20, 21, 33]. This shows that the meaning of utility can be subjective and personal. In particular, for a criminal subject, due to his or her specific biological and psychological characteristics, a desire can be quite deviant from what is commonly considered as the rational norm. For this subject the utility of a certain action a is assessed according to the extent to which it fulfils this personal desire. This shows how utilities are assessed with respect to a subjective measure focusing on a specific desire d, which is affected, or even largely determined by the subject's specific biological and psychological background. According to this perspective, the utility-based decision model was set up as follows:

1. Aspect utility value representations

For any aspect x_i with value s_i, introduce an aspect utility v_i for any possible action a by

has_aspect_utility(a, has_value(x_1, s_1), v_1)
...
has_aspect_utility(a, has_value(x_k, s_k), v_k)

where v_i is based on a closeness measure for each aspect x_i of the considered option a to value s_i, normalised between 0 (least close, minimal utility) and 1 (most close, maximal utility). For example,

has_aspect_utility(fight,
has_value(desire_for_aggressiveness, high), 0.9)

indicates that the action of fighting contributes much to a high value for aggressiveness.

2. Aspect weight factor representations

Introduce weight factors $w_1, \ldots, w_k$ for the different aspects x_i, normalised so that the sum is 1, and introduce relations weight_factor(x_i, w_i) stating that aspect x_i has weight factor w_i.

3. Combination of aspect utilities to option utilities

Combine the option aspect utility values $v_1, \ldots, v_k$ for a given composed desire to an overall option utility taking into account the weight factors $w_1, \ldots, w_k$, according to some combination function $f(v_1, \ldots, v_k, w_1, \ldots, w_k)$.

The combination function in **3.** can be formalised in a number of manners; two common possibilities are:

- Euclidian Distance: $f(v_1, \ldots, v_k, w_1, \ldots, w_k) = \sqrt{(w_1 v_1{}^2 + \cdots + w_k v_k{}^2)}$
- Manhattan Distance: $f(v_1, \ldots, v_k, w_1, \ldots, w_k) = w_1 v_1 + \cdots + w_k v_k$

The LEADSTO property for combination is:

LP41 $\forall$a:ACTION $\forall x_1, \ldots, x_k$:ASPECT $\forall s_1, \ldots, s_k$:SCALE $\forall$ $v_1, \ldots, v_k, w_1, \ldots, w_k$:REAL
belief(has_aspect_utility(a, has_value(x_1, s_1),v_1)) $\wedge \cdots \wedge$
belief(has_aspect_utility(a, has_value(x_k, s_k), v_k)) $\wedge$
weight_factor(x_1, w_1) $\wedge \cdots \wedge$ weight_factor(x_k, w_k) $\twoheadrightarrow_{0,0,1,1}$
belief(has_utility(a, d(has_value(x_1, s_1), ..., has_value(x_k, s_k)), $f(v_1, \ldots, v_k, w_1, \ldots, w_k)$))

Note that the model currently assumes that a value is filled in for each aspect utility and each weight factor. However, in order to deal with missing values, it can easily be extended by defining a default value (for example, 0.5 for the aspect utilities, and 0 for the weight factors) for each of the missing factors.
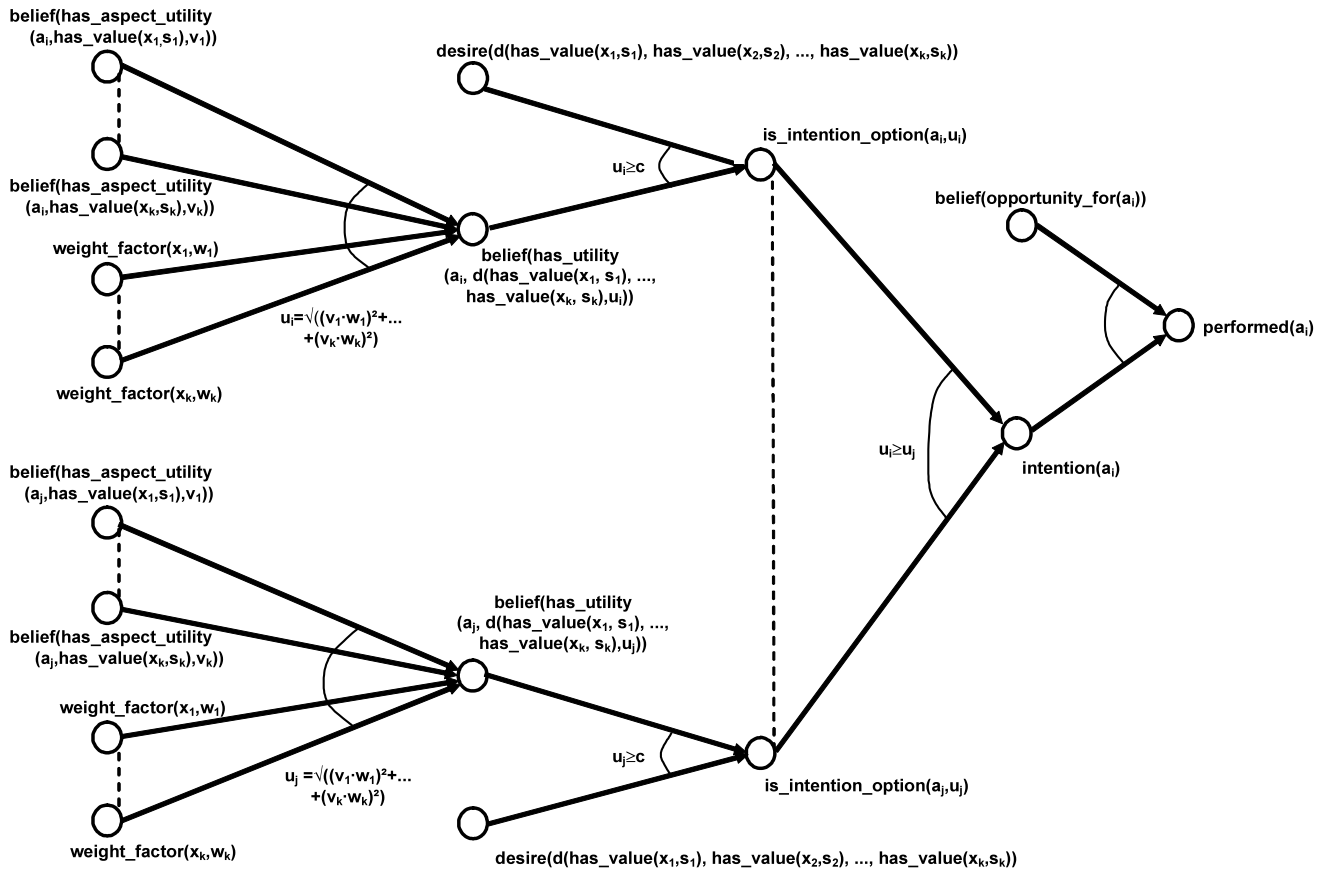
Fig. 2 Utility-based BDI-model

Next, the choice process is formalised. This is done in two steps. First, LP31 is replaced by LP31a, LP31b, and LP31c:

LP31a Desire d combined with the belief that a certain action a will lead to the fulfillment of d with utility u ($\geq$c) will lead to the consideration of a as a possible intention option.
$\forall$d:DESIRE $\forall$a:ACTION $\forall$u:REAL desire(d) $\wedge$
belief(has_utility(a, d, u) $\wedge$ u$\geq$c) $\twoheadrightarrow_{0.2,0.2,1,1}$ is_intention_option(a, u)

Here c is a threshold value, for example 0.5. This is used to generate the options to be considered. To obtain only the intentions with highest utility, as a next phase, the selection process is modelled in two steps by:

LP31b If a1 and a2 are both intention options, but a2 has a higher utility, then a1 is ruled out as an intention option.
$\forall$a1,a2:ACTION $\forall$u1,u2:REAL is_intention_option(a1,u1) $\wedge$
is_intention_option(a2,u2) $\wedge$ u1$<$u2 $\twoheadrightarrow_{0,0,1,1}$
ruled_out_intention_option(a1, u1)

LP31c Eventually, an intention option that is not ruled out is selected as final intention.
$\forall$a:ACTION $\forall$u:REAL is_intention_option(a, u) $\wedge$
not ruled_out_intention_option(a, u) $\twoheadrightarrow_{0,0,1,1}$ intention(a)

The complete utility-based decision model is depicted graphically in Fig. 2. The circles denote state properties, and

the arrows denote dynamic (LEADSTO) properties. Notice that the state properties of the type desire(...) are generated by the model described in the previous section.

Note that, in order to describe a specific decision making scenario with this model, the agent described needs to have some expectancy about possible actions already at the start of the scenario. This expectancy may be triggered by observations (e.g., "I see a potential victim and no guardians, so I consider robbing this person"), or by other internal states (e.g., "I feel like seeking some thrill, so I consider robbing a bank this afternoon"). In the first case, the duration between the decision and the actual performance of the action is rather short, so that it is very likely that an opportunity for the considered action will indeed occur. In the second case, this duration will be longer, and it is possible that no opportunity will occur at all. The model can be used to simulate both types of processes.

6 Example simulation traces

Based on the model shown above, a number of simulation experiments have been performed to test (for some simple scenarios) whether it shows the expected behaviour. In

Table 1 Characteristics of criminal1 and possible assaults[a]

	weight factor (criminal1)	aspect utility (assault1)	aspect utility (assault2)
theory of mind	0.04	low, 0.7	low, 0.9
desire for aggressiveness	0.04	high, 0.3	high, 0.8
desire to act	0.17	high, 0.8	high, 0.8
desire to act safely	0.02	high, 0.7	high, 0.1
desire for actions with strong stimuli	0.17	high, 0.6	high, 0.8
desire for impulsiveness	0.12	medium, 0.5	medium, 0.5
positive emotional attitude towards others	0.02	low, 0.7	low, 0.8
negative emotional attitude towards others	0.04	low, 0.3	low, 0.3
desire for high gain	0.19	high, 0.5	high, 0.8
desire for low loss	0.19	high, 0.8	high, 0.5

[a]This approach assumes that an individual's preferences (i.e., the weight factors), as well as the characteristics of certain actions (i.e., aspect utilities), can be expressed by real numbers. For the presented examples, the chosen numbers are not necessarily claimed to be realistic, and should rather be seen as estimations that were chosen to create interesting scenarios that roughly correspond to reality. All parameter settings were chosen after discussions with domain experts (that were taken from the expert group mentioned in Sect. 2)

this section, two example simulation traces are described in detail. Both traces address the same scenario, but the personal characteristics of the main character differ between the two traces. The first example trace involves a virtual street robber agent (indicated by criminal1) who observes some possible targets, and is deliberating about whether or not to perform an assault (and if so, which assault to perform). For simplicity, we assume that there are two possible assaults to choose from (indicated by assault1 and assault2, respectively). In case of assault1, the agent would steal an old lady's grocery bag, without using extreme violence. In case of assault2, it would steal a young man's brand new laptop. However, since this man seems to be rather strong, it would probably have to use violence to achieve its goal. The characteristics of both assaults, as well as criminal1's individual preferences, are shown in Table 1.

In the second column of the table, the different weight factors assigned to criminal1 can be seen. These weight factors, which add up to 1, show the relative importance of each aspect for the agent. The weight factor for desire to act safely, for example, is 0.02. This means that criminal1 has a low interest in the desire to act safely. The weight factor for desire for actions with strong stimuli is 0.17, which means that he has a high desire for actions with strong stimuli. In the columns to the right of the weight factor, the utility of the different aspects is mentioned (in the third column for assault1 and in the column to the right for assault2). The values describe in how far the aspect is present in this particular assault. For example, has_aspect_utility(assault1, has_value(desire_for_aggressiveness, high), 0.3) shows that assault1 does not contribute much to the high desire for aggressiveness. On the other hand, has_aspect_utility(assault2,

has_value(desire_for_aggressiveness, high), 0.8) shows that assault2 contributes much to the high desire for aggressiveness.

The results of applying the simulation model to this example situation are shown in Fig. 3. Here, time points are on the horizontal axis, whereas the different state properties are on the vertical axis. A box on top of a line indicates that a state property is true at that time point. As shown by this figure, the agent immediately has a certain desire, represented as d1. Note that this stands for a complex desire represented as:

```
d(has_value(theory_of_min,low),
  has_value(desire_for_aggressiveness,high),
  has_value(desire_to_act,high), has_value(desire_to_act_safely,high),
  has_value(desire_for_actions_with_strong_stimuli,high),
  has_value(desire_for_impulsiveness,medium),
  has_value(positive_emotional_attitude_towards_others,low),
  has_value(negative_emotional_attitude_towards_others,low),
  has_value(desire_for_high_gain,high),
  has_value(desire_for_low_loss,high))
```

(which was not shown in the picture, for obvious reasons). This desire was generated by a complex process, involving a combination of biological and psychological factors. For presentation purposes, this part of the trace is not shown here either. However, more detailed simulation traces that include such processes are shown in Appendix A in [5].

Based on the desire as described above, agent criminal1 then starts assessing the utilities of the two possible assaults (see the predicates belief(has_utility(...)) at time point 1), based on the aspect utilities and weight factors of these assaults. The action of stealing the young man's laptop (assault2) is assessed with value 0.678723, whereas the action of robbing the old lady's groceries (assault1) has value
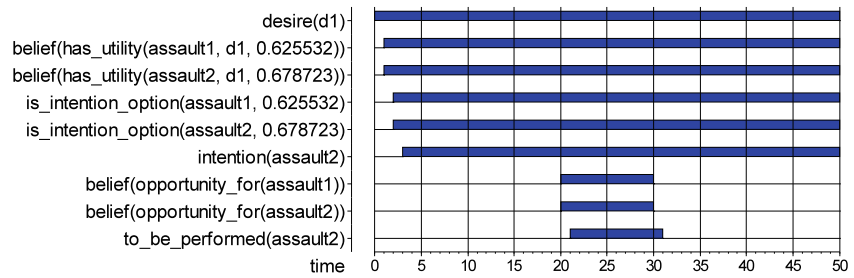
Fig. 3 Example simulation trace for criminal 1

Table 2 Characteristics of criminal 2 and possible assaults

	weight factor (criminal2)	aspect utility (assault1)	aspect utility (assault2)
theory of mind	0.04	low, 0.7	low, 0.9
desire for aggressiveness	0.04	high, 0.3	high, 0.8
desire to act	0.17	high, 0.8	high, 0.8
desire to act safely	0.04	high, 0.7	high, 0.1
desire for actions with strong stimuli	0.15	high, 0.6	high, 0.8
desire for impulsiveness	0.12	medium, 0.5	medium, 0.5
positive emotional attitude towards others	0.02	low, 0.7	low, 0.8
negative emotional attitude towards others	0.04	low, 0.3	low, 0.3
desire for high gain	0.11	high, 0.5	high, 0.8
desire for low loss	0.28	high, 0.8	high, 0.5

0.625532. Since both has_utility-values are higher than 0.5, both actions become possible intentions (see time point 2). Next, the agent chooses the one with the highest utility, which leads to the intention to perform assault2 at time point 3. Later, when an opportunity for assault2 arises (time point 20), this assault is indeed performed (time point 21).

In a second example trace another street robber agent (criminal 2) observes the same possible targets as criminal 1, and is also deliberating about whether or not to perform an assault (and if so, which assault to perform). Also in this case, we assume that there are two possible assaults to choose from: assault1 (stealing an old lady's grocery bag without using extreme violence) and assault2 (stealing a young man's brand new laptop potentially with the use of violence). The characteristics of both assaults are the same as in the previous case study. However, the preferences of the two criminals are somewhat different. The individual preferences of criminal2 are shown in Table 2.

In the first column of the table, the different weight factors assigned to criminal2 can be seen. These weight factors, which add up to 1, show the relative importance of each aspect for the agent. Compared to criminal1, criminal2 has a different weight factor assigned to the desires to act safely, for actions with strong stimuli, for high gain and for low loss. The weight factor for desire for high gain, for example, is 0.11. This means that criminal2 has a high interest in the desire for high gain. However, the weight factor for

desire for low loss 0.28, which means that the agent has a very high desire for low loss. For this criminal, low loss is much more important than high gain. In the columns to the right of the weight factor, the utility of the different aspects is mentioned.

The results of applying the simulation model to this example situation are shown in Fig. 4. As shown by this figure, the agent immediately has a certain desire, represented as d2, which stands for a more complex desire represented as:

```
d(has_value(theory_of_min,low),
   has_value(desire_for_aggressiveness,high),
   has_value(desire_to_act,high), has_value(desire_to_act_safely,high),
   has_value(desire_for_actions_with_strong_stimuli,high),
   has_value(desire_for_impulsiveness,medium),
   has_value(positive_emotional_attitude_towards_others,low),
   has_value(negative_emotional_attitude_towards_others,low),
   has_value(desire_for_high_gain,high),
   has_value(desire_for_low_loss,high))
```

as can also be seen in the last two columns of Table 2. Based on the desire as described above, criminal2 then starts assessing the utilities of the two possible assaults (see the predicates belief(has_utility(...)) at time point 1), based on the aspect utilities and weight factors of these assaults. The action of stealing the old lady's groceries (assault1) is assessed with value 0.652475, whereas the action of robbing the young man's laptop (assault2) has value 0.637624. Since

Fig. 4 Example simulation trace for criminal 2
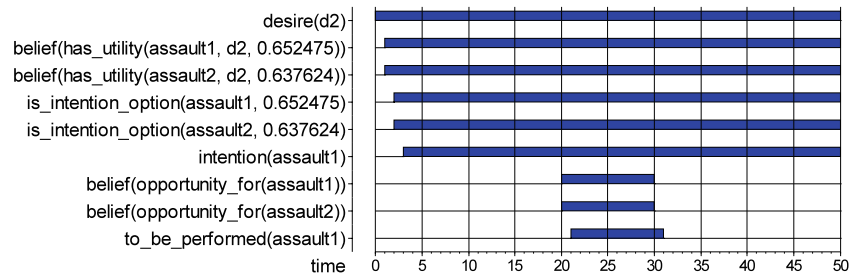
both has_utility-values are higher than 0.5, both actions become possible intentions (see time point 2). Next, the agent chooses the one with the highest utility, which leads to the intention to perform assault1 at time point 3. Later, when an opportunity for assault1 arises (time point 30), this assault is indeed performed (time point 31).

As illustrated by the traces in Figs. 3 and 4, the simulation experiments have indicated that the presented model successfully integrates personal biological and psychological aspects within the decision making process, which eventually leads to the selection of actions that correspond to the desires of the individual. Although the presented scenario was kept simple, these experiments illustrate how the model allows agents with different personal characteristics to make different decisions in the same situation (criminal1 decided to rob the old lady, whereas criminal2 decided to steal the young man's laptop).

In addition to the simulations described above, the model has been used to generate various other simulation traces under different parameter settings. Such simulation experiments enabled the modellers to investigate (at an abstract level) whether the developed model shows satisfactory behaviour (cf. prototyping in Software Engineering). The obvious next step is to implement the model within concrete real-world applications. Due to the intuitive (causal relationship-based) format of the LEADSTO language, and the fact that it is independent of a particular implementation language, this step is relatively straightforward. For earlier work where agents models written in LEADSTO have been converted to virtual environment applications, see [8, 9].

7 User study

To evaluate how humans perceive the developed agent (with our criminal decision making model) we performed an experiment. We used the scenario introduced earlier, about a street robber who has the choice between different assaults. To be able to assess the impact of the different parts of the criminal decision making model, we applied some variations of the model to the scenario.

A. The criminal decision making model with the personality traits of criminal 1.

B. The criminal decision making model with the personality traits of criminal 2.
C. Standard BDI model without any personality traits.
D. The criminal decision making model with the personality traits of criminal 1, where we deliberately implemented the decision making component incorrect. This means that the model suggests those actions that are least satisfactory to the agent's desire as best possibilities.
E. The criminal decision making model with the personality traits of criminal 2, where we deliberately implemented the action generation component incorrect. This means that the model selects actions that do not match the agent's intentions.

Variant D and E have been added to check whether virtual agent believability is already enhance by a criminal decision making model per se, or whether a realistic model is needed.

After generating simulation traces on the basis of these 5 models, for each of the atomic state properties that occur in the model, a mapping has been create to a text fragment. For example, the state property

$$is\_intention\_option(X,Y) \text{ (where } Y \text{ is high)}$$

corresponds to the text fragment

> The agent considers X a good option.

Similarly, the state property

$$belief(opportunity\_for(X))$$

corresponds to

> The agent believes that there is an opportunity for X.

Using these mappings and a specific conversion program that has been written (see [4]), the LEADSTO simulation traces are automatically translated into virtual storylines in textual format. An example of such a generated storyline is shown in Fig. 5.

Twenty persons participated in the experiment. The age of the participants ranged from 23 to 58 with a mean age of 32 and a standard deviation of 9. Among the participants 13 were male and 7 were female. Each participant was asked to read the five different stories (there were different ver-
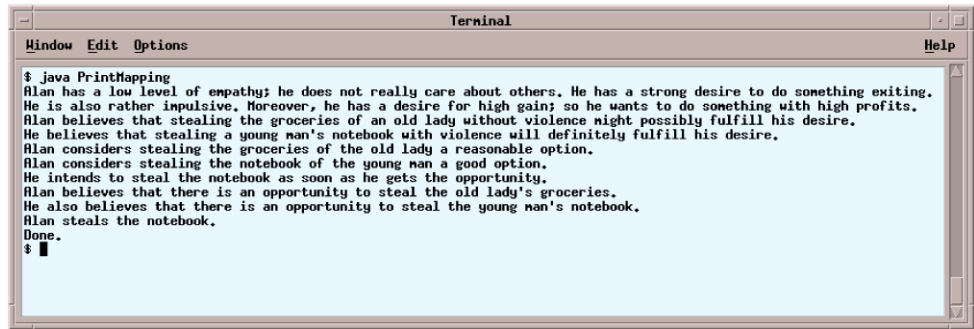
Fig. 5 Example of a generated storyline

```
                              Terminal
 Window  Edit  Options                                          Help

 $ java PrintMapping
 Alan has a low level of empathy; he does not really care about others. He has a strong desire to do something exiting.
 He is also rather impulsive. Moreover, he has a desire for high gain; so he wants to do something with high profits.
 Alan believes that stealing the groceries of an old lady without violence might possibly fulfill his desire.
 He believes that stealing a young man's notebook with violence will definitely fulfill his desire.
 Alan considers stealing the groceries of the old lady a reasonable option.
 Alan considers stealing the notebook of the young man a good option.
 He intends to steal the notebook as soon as he gets the opportunity.
 Alan believes that there is an opportunity to steal the old lady's groceries.
 He also believes that there is an opportunity to steal the young man's notebook.
 Alan steals the notebook.
 Done.
 $
```
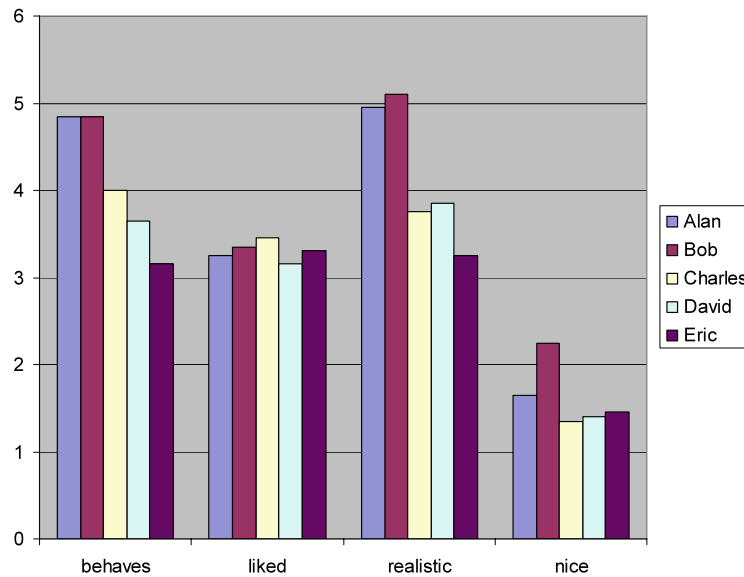
Fig. 6 Statistical results of the experiment

Table 3 Detailed statistical results

Statement	A-B	A-C	A-D	A-E	B-C	B-D	B-E	C-D	C-E	D-E	Overall
Behaves	n.s.	n.s.	0.023	0.002	n.s.	0.024	0.002	n.s.	n.s.	n.s.	0.003
Liked	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.	n.s.
Realistic	n.s.	0.013	0.02	0.000	0.005	0.008	0.000	n.s.	n.s.	n.s.	0.000
Nice	n.s.	n.s.	n.s.	n.s.	0.009	0.009	0.014	n.s.	n.s.	n.s.	0.008

sions of the experiment, each with a different order of the stories to avoid ordering problems). After reading each scenario the participant filled out a questionnaire. This questionnaire consisted of 12 statements (see the Appendix). The participants were asked to award a measurement of agreement on these 12 statements, about how they perceived the agent. A gradual seven-point scale was used, with the following meaning: 1 = 'I strongly disagree', 2 = 'I disagree', 3 = 'I weakly disagree', 4 = 'neutral', 5 = 'I weakly agree', 6 = 'I agree', 7 = 'I strongly agree'.

7.1 Results

To analyse the results of the experiment, an ANOVA has been applied on the answers to the statements of the ques-

tionnaire. In this section, the most relevant statements are addressed. These statements are: 'The agent behaves like a real person', 'I liked reading this story', 'This is a realistic story', and 'The agent seems to be a nice guy'.

Recall that there were five variants of the scenario: variant A-E introduced earlier. The results are presented in Table 3 and Fig. 6. The vertical axis in Fig. 6 corresponds to the scale explained above. The first column of the table indicates the statement, and the other columns indicate pair wise comparisons between the different variants. The "total" column shows whether the comparison yielded a significant result or not. The cells indicate the p-value of the ANOVA (in case the difference was significant), or 'n.s.' (in case the difference was not significant). For example, the third cell of the first row states '0.023', which means that, according to

the participants, variant A behaved significantly more like a real person than variant D.

These results show that there is no significant difference between how much the different agents are liked by the participants (second row). There is a significant difference between agent A and B on the one hand and agent D and E on the other hand when it comes to their behaviour (first row). Agents A and B are the agents that behave according to the criminal decision making model. Agents D and E on the other hand, used an incorrect version of the decision making model. The participants clearly feel that agents A and B behave like real persons and agents D and E do not. There is also a significant difference between A and B on the one hand and C, D and E on the other hand when it comes to how realistic the stories are (third row). The participants agree that A and B are more realistic stories than C, D and E. Further, there is a significant difference between how nice the participants find agent B compared to agent C, D and E. This is an interesting finding. In story B the agent steals a bag of groceries without any violence, but in addition there is some background knowledge on his personal characteristics which he uses in his decision making. In story D, the agent also steals the bag of groceries, but has different personal characteristics, which do not match his behaviour. Surprisingly, this yielded the participants to perceive this agent as less friendly.

8 Discussion

In this paper, a model for decision making in virtual agents is presented, which combines a BDI-agent model with a model for multi-attribute decision making. It enables a choice between different options for actions fulfilling a complex desire, according to the Rational Choice Theory in Criminology. The resulting agent model combines qualitative, logical aspects of a BDI-model with quantitative, numerical aspects of utility theory. The model is illustrated by applying it to criminal decision making (and more specifically, focussing on certain types of violent criminals). To this end, [6] criminal agent model was used for desire generation. To study the behaviour of the model, it has been applied to a case of street robbery, for which various scenarios have been simulated. In addition, a user study has been performed, in which the generated simulation traces were converted into virtual storylines, which were read by 20 participants. The study confirmed our hypothesis that the model enhances believability of virtual agents. The participants appreciated those stories that were generated on the basis of the decision making better than baseline stories that were generated on the basis of a traditional BDI model, or an 'incongruous' variant of the decision making model.

Despite these results, a complete external validation of the model remains a nontrivial issue. At least, the present paper has indicated that it is possible to integrate biological and psychological factors with rational factors within one model. Moreover, the model indeed shows the behaviour of different types of criminals as described in literature such as [13, 27, 34], and the user study has confirmed that this leads to a higher believability. In this sense the model has been validated positively. However, notice that this is a relative validation, only with respect to the literature on criminal decision making, and the expectations by humans. In cases that the available knowledge about the behaviour and biological and psychological functioning of such criminal types is extended, the model can be validated accordingly and when needed improved. The modelling approach as put forward supports such an incremental development and improvement. The simulation model has been specified in a conceptual, not implementation-dependent manner, and is easy to maintain. In this sense the approach anticipates further development of the area of criminal behaviour.

Concerning related work, in the literature in cognitive science, more and more authors propose *dual process theories*, which claim that cognition can be divided into two distinct systems: a low-level, emotional and unconscious system, and a high-level, evolutionary recent, conscious system, see, e.g., [14, 40, 41]. As a result, most contemporary cognitive modelling architectures (e.g., ACT-R [1], Soar [23], and, more recently, Clarion [40, 41]) have evolved under the influence of these theories. For example, although ACT-R is traditionally classified as a 'symbolic' architecture (mainly aimed at representing logical, qualitative aspects), it has more recently been extended with 'subsymbolic' mechanisms to represent quantitative, numerical aspects as well. Currently, the subsymbolic part of ACT-R is represented by a large set of parallel processes that can be summarised by a number of mathematical equations, whereas its symbolic part is fulfilled by a production system. Here, the subsymbolic equations control many of the symbolic processes. For instance, if multiple production rules in ACT-R's symbolic part are candidates to be executed, a subsymbolic utility equation may estimate the relative cost and benefit associated with each rule and select the rule with the highest utility for execution. A similar development holds for Soar, although the distinction between symbolic and subsymbolic is a bit less evident here. Clarion, instead, has always been classified as a 'hybrid' cognitive architecture.

Also various cognitive theories about decision making itself are inspired by the idea to distinguish two different types of processes. For example, the recognition-primed decision (RPD) theory (see [22]) for the theory, and [42] for a computational model) describes decision making as a process that consists of an 'intuitive' part in which the current situation is matched against patterns learned from experiences and a 'conscious, deliberate' part in which the consequences of actions are evaluated. Recently, the RPD theory has been

applied in the area of multi-agent systems as well, see e.g. [15, 16, 28], although these applications have not addressed virtual storytelling as yet.

At first sight, our proposed model seems to show significant similarities with such theories: our model to determine desires has characteristics of an unconscious, low-level system, whereas the model for utility-based decision making resembles a conscious, high-level system. Future work will explore whether a more precise mapping can be made between the concepts introduced in our combined model and the concepts typically used in dual process theory. One of the differences with existing 'dual process' approaches, seems to be the fact that our model takes biological factors into account. Another difference is that the cognitive architectures mentioned above are more generic, whether the model as presented in this paper is specialised towards modelling criminal behaviour. Finally, the original research goals underlying both methodologies are somewhat different. Cognitive architectures have (at least traditionally) as main purpose to study cognitive processes, whether the model proposed here aims at developing believable agents.

Also when comparing our model with numerical approaches to model decision making, such as neural or connectionist networks [37], a difference in perspective can be observed. The main aim of the latter class of methods is to build intelligent behaviour (e.g., to solve certain optimisation problems), whereas our approach was designed to enhance realism (and not necessarily intelligence). This difference also (partly) explains why the approach taken in this article was to use build an explicit model using existing knowledge from the literature, instead of taking a 'black box' approach, where it is difficult to understand the actual behaviour of the model. On the other hand, we acknowledge that the presented model does not incorporate a mechanism to improve itself in an automated manner, as is done in machine learning. Nevertheless, the model can be improved manually in an iterative manner, which is mainly done by 1) changing the settings of the parameters involved, or by 2) adding or removing factors to the subconscious model. Input for the first type of adaptation can be derived from user studies such as described in Sect. 7, whereas the second type of adaptation can be performed based on new insights in the expert literature on criminal decision making (as described earlier).

In addition, it is interesting to explore how our model relates to theories in which affective factors just 'bypass' decision making, such as in [12, 24]. In recent years, such theories have formed a source of inspiration for the development of a number of agent-based models for reasoning and (affective) decision making (which often also combine rational and non-rational aspects) within Artificial Intelligence, such as in [19, 26, 30, 39]. A difference with the presented model is that most of these models explicitly focus on the integration of *emotions* with rational behaviour, whereas our model

tries to integrate rational behaviour with personal biological and psychological factors *in general* (including emotions, but also notions like aggressiveness and impulsiveness). Finally, another interesting direction for future work will be to investigate whether the presented model can easily be implemented in standard BDI-based agent modelling frameworks, such as AgentSpeak [35] or Jadex [31]. Since the basis of our model is very similar to those models, this seems like a promising direction.

Appendix: Storylines and questionnaire

Story 1: Alan

Alan has a low level of empathy; he does not really care about others. He has a strong desire to do something exiting. He is also rather impulsive and aggressive. Moreover, he has a desire for high gain; so he wants to do something with high profits. Alan believes that stealing the groceries of an old lady without violence might possibly fulfill his desire. Further he believes that stealing a young man's notebook with violence will definitely fulfill his desire. Alan considers stealing the groceries of an old lady a reasonable option. Alan considers stealing the notebook of a young man a good option. He intends to steal the notebook as soon as he gets the opportunity. Alan believes that there is an opportunity to steal an old lady's groceries. He also believes that there is an opportunity to steal a young man's notebook. Alan steals the notebook.

Story 2: Bob

Bob has a low level of empathy; he does not really care about others. He has a desire to do something exiting, but if possible, he wants to do something that is safe. He is also rather impulsive. He has a desire for high gain, but he has an even higher desire for low loss. So, he wants to do something with high profit, but thinks that low risks are even more important. Bob believes that stealing the groceries of an old lady without violence will definitely fulfill his desire. Bob believes that stealing a young man's notebook with violence

might possibly fulfill his desire. Bob considers stealing the groceries of an old lady to be a good option. Bob considers stealing the notebook of a young man to be a reasonable option. He intends to steal the groceries as soon as he gets the opportunity. Bob believes that there is an opportunity to steal a young man's notebook and that there is an opportunity to steal an old lady's groceries. Bob steals the groceries.

Story 3: Charles

Charles has a desire to assault someone. He believes that robbing the groceries from an old lady will fulfill his desire. He also believes that stealing the notebook from a young man will fulfill his desire. Charles has the intention to rob an old lady. He also has the intention to steal a notebook. Charles believes that there is an opportunity to steal a young man's notebook. He steals the notebook. He also believes that there is an opportunity to steal an old lady's groceries. Charles steals the groceries.

Story 4: David

David has a low level of empathy; he does not really care about others. He has a strong desire to do something exiting. He is also rather impulsive and aggressive. Moreover, he has a desire for high gain; so he wants to do something with high profits. David believes that stealing the groceries of an old lady without violence will definitely fulfill his desire. Further he believes that stealing a young man's notebook with violence might possibly fulfill his desire. David considers stealing the groceries of the old lady to be a good option. David considers stealing the notebook of a young man to be a reasonable option. He intends to steal the groceries as soon as he gets the opportunity. David believes that there is an opportunity to steal an old lady's groceries. He also believes that there is an opportunity to steal a young man's notebook. David steals the groceries.

Story 5: Eric

Eric has a low level of empathy; he does not really care about others. He has a desire to do something exiting, but if possible, he wants to do something that is safe. He is also rather impulsive. He has a desire for high gain, but he has an even higher desire for low loss. So, he wants to do something with high profit, but thinks that low risks are even more important. Eric believes that stealing the groceries of an old lady without violence will definitely fulfill his desire. Eric believes that stealing a young man's notebook with violence might possibly fulfill his desire. Eric considers stealing the groceries of an old lady to be a good option. He considers stealing the notebook of a young man to be a reasonable option. He intends to steal the groceries as soon as he gets the

opportunity. Eric believes that there is an opportunity to steal a young man's notebook and that there is an opportunity to steal an old lady's groceries. Eric steals the notebook.

Questionnaire

Name:
Gender:
Age:
Occupation:

Please fill out the following questionnaire
(1 = totally disagree, 4 = neutral, 7 = totally agree)

I expected Alan to steal the notebook	1	2	3	4	5	6	7
Alan behaves like a real person	1	2	3	4	5	6	7
I understand Alan's behaviour	1	2	3	4	5	6	7
Alan is an aggressive person	1	2	3	4	5	6	7
I liked reading this story	1	2	3	4	5	6	7
This is an interesting story	1	2	3	4	5	6	7
This is a realistic story	1	2	3	4	5	6	7
I have sympathy for Alan	1	2	3	4	5	6	7
Alan seems to be a nice guy	1	2	3	4	5	6	7
Alan thinks before he acts	1	2	3	4	5	6	7
Alan behaves rationally	1	2	3	4	5	6	7
Alan reasons rationally	1	2	3	4	5	6	7

References

1. Anderson JR, Lebiere C (1998) The atomic components of thoughts. Erlbaum, Mahwah
2. Baron-Cohen S (1995) Mindblindness. MIT Press, Cambridge
3. Bartol CR (2002) Criminal behavior: a psychosocial approach, 6th edn. Prentice Hall, New York
4. Bosse T, Gerritsen C, Treur J (2007) Case analysis of criminal behaviour. In: Proceedings of the 20th international conference on industrial, engineering and other applications of applied intelligent systems, IEA/AIE'07. Lecture notes in computer science, vol 4570. Springer, Berlin, pp 621–632
5. Bosse T, Gerritsen C, Treur J (2007) Cognitive and social simulation of criminal behaviour: the intermittent explosive disorder case. In: Proceedings of the sixth international joint conference on autonomous agents and multi-agent systems, AAMAS'07. ACM Press, New York, pp 367–374
6. Bosse T, Gerritsen C, Treur J (2007) Integration of biological, psychological and social aspects in agent-based simulation of a violent psychopath. In: Shi Y, van Albada GD, Dongarra J (eds) Proceedings of the seventh international conference on computational science, ICCS'07. Lecture notes in computer science, vol 4488. Springer, Berlin, pp 888–895

7. Bosse T, Jonker CM, van der Meij L, Treur J (2007) A language and environment for analysis of dynamics by SimulaTiOn. J AI Tools 16:435–464

8. Bosse T, de Lange FPJ (2008) Development of virtual agents with a theory of emotion regulation. In: Jain L, Gini M, Faltings BB, Terano T, Zhang C, Cercone N, Cao L (eds) Proceedings of the eighth IEEE/WIC/ACM international conference on intelligent agent technology, IAT'08. IEEE Computer Society Press, Los Alamitos, pp 461–468

9. Bosse T, Pontier M, Siddiqui GF, Treur J (2007) Incorporating emotion regulation into virtual stories. In: Pelachaud C, Martin JC, Andre E, Chollet G, Karpouzis K, Pele D (eds) Proceedings of the seventh international conference on intelligent virtual agents, IVA'07. Lecture notes in artificial intelligence, vol 4722. Springer, Berlin, pp 339–347

10. Cavazza M, Charles F, Mead S (2003) Interacting with virtual characters in interactive storytelling. In: Alonso E, Kudenko D, Kazakov D (eds) Adaptive agents and multi-agent systems. Lecture notes in artificial intelligence, vol 2636. Springer, Berlin, pp 318–325

11. Cornish DB, Clarke RV (1986) The reasoning criminal: rational choice perspectives on offending. Springer, Berlin

12. Damasio AR (1994) Descartes' error: emotion, reason, and the human brain. Putnam, New York

13. Delfos MF (2004) Children and behavioural problems: anxiety, aggression, depression and ADHD; a biopsychological model with guidelines for diagnostics and treatment. Harcourt, Amsterdam

14. Evans JSBT (2003) In two minds: dual-process accounts of reasoning. Trends Cogn Sci 7:454–459

15. Fan X, Sun S, McNeese M, Yen J (2005) Extending the recognition-primed decision model to support human-agent collaboration. In: Proceedings of the fourth international joint conference on autonomous agents and multiagent systems, AAMAS'05. ACM Press, New York, pp 945–952

16. Fan X, Sun B, Sun S, McNeese M, Yen J (2006) RPD-enabled agents teaming with humans for multi-context decision making. In: Proceedings of the fifth international joint conference on autonomous agents and multiagent systems, AAMAS'06. ACM Press, New York, pp 34–41

17. Georgeff MP, Lansky AL (1987) Reactive reasoning and planning. In: Proc of the sixth national conf on artificial intelligence, AAAI'87, Menlo Park, CA. American Association for Artificial Intelligence, Menlo Park, pp 677–682

18. Hoorn JF, Konijn EA, Van der Veer GC (2003) Virtual reality: do not augment realism, augment relevance. Upgrade—Human-Comput Interact Overcoming Barriers 4(1):18–26

19. Jiang H, Vidal JM, Huhns MN (2007) EBDI: an architecture for emotional agents. In: Proc of the 6th int joint conf on autonomous agents and multi-agent systems, AAMAS'07. ACM Press, New York, pp 38–40

20. Jonker CM, Treur J (2001) An agent architecture for multi-attribute negotiation. In: Nebel B (ed) Proc of the 17th international joint conference on AI, IJCAI'01. Morgan Kaufman, San Mateo, pp 1195–1201

21. Keeney R, Raiffa H (1976) Decisions with multiple objectives: preferences and value tradeoffs. Wiley, New York

22. Klein G (1989) Recognition-primed decisions. Adv Man-Mach Syst Res 5:47–92

23. Laird JE, Newell A, Rosenbloom PS (1987) Soar: an architecture for general intelligence. Artif Intell 33(1):1–64

24. Loewenstein GF, Weber EU, Hsee CK, Welch ES (2001) Risk as feelings. Psychological Bull 127:267–286

25. Maes P, Guttman RH, Moukas AG (1999) Agents that buy and sell. Commun ACM 42(3):81–91

26. Marsella S, Gratch J (2009) EMA: a model of emotional dynamics. Cogn Syst Res 10(1):70–90

27. Moir A, Jessel D (1995) A mind to crime: the controversial link between the mind and criminal behaviour. Michael Joseph, Penguin, London

28. Norling E, Sonenberg L, Ronnquist R (2000) Enhancing multi-agent based simulation with human-like decision making strategies. In: Moss S, Davidsson P (eds) Proceedings of the second international workshop on multi-agent based simulation, pp 214–228

29. Pelachaud C, Martin JC, Andre E, Chollet G, Karpouzis K, Pele D (eds) (2007) Intelligent virtual agents. In: Proceedings of the 7th international conference on intelligent virtual agents, IVA'07. LNAI, vol 4722. Springer, Berlin

30. Pereira D, Oliveira E, Moreira N (2006) Modelling emotional BDI agents. In: ECAI workshop on formal approaches to multi-agent systems (FAMAS 2006), Riva del Garda, Italy

31. Pokahr A, Braubach L, Lamersdorf W (2005) Jadex: a BDI reasoning engine. In: Bordini RH, Dastani M, Dix J, El Fallah Seghrouchni A (eds) Multi-agent programming: languages, platforms and applications. Springer, Berlin, pp 149–174

32. Prendinger H, Lester J, Ishizuka M. (eds) (2008) Intelligent virtual agents. In: Proceedings of the 8th international conference on intelligent virtual agents, IVA'08. LNAI, vol 5208. Springer, Berlin

33. Raiffa H (1982) The art and science of negotiation. Harvard University Press, Cambridge

34. Raine A (1993) The psychopathology of crime: criminal behaviors as a clinical disorder. Guilford Publications, New York

35. Rao AS (1996) AgentSpeak(L): BDI agents speak out in a logical computable language. In: de Velde WV, Perrame J (eds) Agents Breaking away: proceedings of the seventh European workshop on modelling autonomous agents in a multi-agent world (MAAMAW'96). LNAI, vol 1038. Springer, Berlin, pp 42–55

36. Rao AS, Georgeff MP (1991) Modelling rational agents within a BDI-architecture. In: Allen J et al (eds) Proc of the 2nd int conference on principles of knowledge representation and reasoning, (KR'91). Morgan Kaufmann, San Mateo, pp 473–484

37. Read SJ, Miller LC (eds) (1998) Connectionist models of social reasoning and social behavior. Erlbaum, Mahwah

38. Schaub H (1997) Decision making in complex situations: cognitive and motivational limitations. In: Flin R, Salas E, Strub ME, Martin L (eds) Decision making under stress. Emerging themes and applications. Ashgate, Aldershot, pp 291–300

39. Silverman BG (2001) More realistic human behavior models for agents in virtual worlds: emotion, stress, and value ontologies. Technical Report, Philadelphia, PA, University of Penn/ACASA

40. Sun R (2002) Duality of the mind. Lawrence Erlbaum Associates, Mahwah

41. Sun R (2006) The clarion cognitive architecture: extending cognitive modeling to social simulation. In: Sun R (ed) Cognition and multi-agent interaction. Cambridge University Press, New York, pp 79–99

42. Warwick W, McIlwaine S, Hutton R, McDermott P (2001) Developing computational models of recognition-primed decision making. In: 10th conference on computer generated forces and behavioral representation, Norfolk, VA, 14–17 May 2001, pp 323–331

Tibor Bosse is an Assistant Professor in Artificial Intelligence at Vrije Universiteit Amsterdam. He obtained his Ph.D. in Artificial Intelligence in 2005, graduating on the topic 'Analysis of the Dynamics of Cognitive Processes'. Since June 2006 he is working as Assistant Professor in the Agent Systems Research Group of the Vrije Universiteit. Since 2002, his research has been situated in the intersection of Artificial Intelligence, Cognitive Science, and Ambient Intelligence. His main research interest is computational modeling of human-directed (e.g., cognitive, biological, social) processes, both for theoretical and practical purposes. Moreover, he is a co-organizer of the series of international workshops on Human Aspects in Ambient Intelligence (HAI).

Charlotte Gerritsen is a Ph.D. student in Artificial Intelligence at Vrije Universiteit Amsterdam. Her main research interest addresses computational modeling of criminal behavior. Within Criminology, the analysis of criminal behavior is a central issue. Such an analysis involves different types of aspects, such as biological, psychological and social aspects and their mutual interactions. Usually such analyses are made by criminologist researchers or practitioners, without using any formalisation or computer support. In recent years much progress has been made in biological, cognitive and social complex dynamical systems modelling within areas such as Artificial Intelligence. The methods developed in these areas, when integrated, open the perspective to address the analysis of criminal behavior in more exact, formalised and computer supported manners. Thus, the way is opened to a more solid basis and computer support for simulation and analysis. Charlotte's project explores this potential. It identifies on the one hand useful knowledge from the literature in Criminology and the different disciplines underlying it, and on the other hand a number of dedicated modelling techniques from areas such as Artificial Intelligence and Cognitive Science. By combining these, an integrated computer supported method to criminal behavior analysis will be created.

Jan Treur has a full professorship in Artificial Intelligence at Vrije Universiteit Amsterdam (VUA) since 1990. He is heading the department of Artificial Intelligence consisting of about 45 researchers. He is an internationally well-recognized expert in agent technology, cognitive modeling and knowledge engineering. He is or has been member of the programme committee of many of the main conferences and workshops and many journals in these areas. His extensive list of publications (e.g., see www.cs.vu.nl/~treur) covers major scientific publication media in Artificial Intelligence and Cognitive Science, including the top level conferences and journals. Some of his recent involvements are organizing and charing the series of international workshops on Human Aspects in Ambient Intelligence (HAI). Moreover, he initiated and designed a strongly multidisciplinary Bachelor study program at VUA in Human Ambience, combining subjects from Artificial Intelligence, Computer Science, Psychology and Biomedical Sciences.