A Model for Criminal Decision Making based on Hypothetical Reasoning about the Future

Tibor Bosse and Charlotte Gerritsen

Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081, NL-1081 HV, Amsterdam, The Netherlands
Email: {tbosse, cg}@few.vu.nl
URL: http://www.few.vu.nl/~{tbosse, cg}

Abstract. This paper presents an agent-based model for decision making, which integrates personal biological and psychological aspects with rational utility-based reasoning. The model takes a BDI-based approach, where generation of desires is based on the personal characteristics, and generation of intentions is based on the rational reasoning. Moreover, a hypothetical reasoning mechanism is exploited to derive knowledge that connects certain actions to desires. The model has been implemented in the LEADSTO environment, and has been applied in a case study in the domain of criminal behaviour. Simulation experiments pointed out that the model enables agents to reason effectively about the consequences of their actions, which helps them to make the decisions that best satisfy their desires.

1 Introduction

Decision making is a complex process, which has received a lot of attention within various disciplines, ranging from psychology e.g., [1, 11, 16] and economics [17, 26] to computer science and artificial intelligence [12, 28]. Both for humans and artificial agents, the question 'which factors influence a certain decision' is important and nontrivial. A known problem encountered by modellers of decision making is that this process is determined partly by *rational* means-end reasoning, and partly by subjective personal *biological* and *psychological* aspects (including, for example, a person's motivational and emotional state), see, e.g., [9, 10, 20]. 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 hit someone but may decide rationally not to do this because this may have negative consequences. However, if the biological desire is too strong, the person may decide to hit nevertheless. Thus, humans exploit some mechanism that enables them to make decisions in situations where both rational and biological/psychological factors play a role.

This paper is part of a larger project, which has as main objective to develop a formal agent-based model of such decision making mechanisms. In principle, a generic approach will be taken, i.e., enabling modellers to formalise decision making in any arbitrary domain. More specifically, however, the current paper will focus on 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 personal biological background, or is the result of a rational, calculated choice; e.g., [8, 23]. The current paper will show how the two viewpoints can be integrated, thereby creating a behavioural model for a "criminal agent". Such a model can be useful from two perspectives. First, from a theoretical point of view, it may be used to get more insight in the process of criminal decision making itself. Second, from an application point of view, it may be used to develop *virtual agents* cf. [25] that are able to show criminal behaviour. Typical applications in which such criminal agents could play a role are (serious) games and computer-generated virtual stories e.g., [7].

As a starting point, the model described in [3] will be taken. This model addresses criminal action generation based on beliefs, desires and intentions (BDI), where generation of *desires* is based on biological and psychological aspects (such as high levels of testosterone or serotonin), and generation of *intentions* is based on rational, utility-based multi-criteria decision making cf. [17, 26]. According to this model, in case an agent has, for example, a high level of testosterone, this agent will generate a high desire for aggressive actions. When combining this desire with the belief that a certain action A will satisfy the desire, a state is generated in which A is considered a possible action. Next, this state is compared with other alternatives during a multi-criteria decision making process. However, one drawback of the model in [3] is that the knowledge that 'a certain action A satisfies a desire D' is assumed given. In contrast, in real world situations, humans do not always have such knowledge directly available. Instead, they often need to perform some kind of *hypothetical reasoning* [15, 30], in order to determine what will be the consequence of a certain (potential) action. Such reasoning may involve some derivation of possible effects in the world of a certain action, but also of effects on other persons (social cognition). For example, a criminal that is deliberating about whether or not to rob an old lady may take into account certain positive effects on the world (e.g., he will gain some money) and on other persons (e.g., his friends will admire him), as well as negative effects on the world (e.g., he might get caught) and on other persons (e.g., his parents will be disappointed). This paper extends the model from [3] with a model to determine such potential consequences of hypothetical actions.

Section 2 provides an overview of the literature used as a basis for the proposed model. In Section 3, the decision making model is presented, and Section 4 presents some simulation results. Section 5 provides a discussion about the work.

2 Hypothetical Reasoning, Social Cognition, and Rational Choice

Hypothetical reasoning[1] (also sometimes called 'what-if reasoning', or 'reasoning by assumption') is a practical reasoning pattern that often occurs within everyday life [15, 30]. The circumstances in which the pattern may be applied vary from diagnostic problems to cases where persons deliberate about which action to perform. An

[1] Note that this paper only addresses hypothetical reasoning about the future, in a way that is similar to temporal projection [13]. As such, it is different from hypothetical reasoning about the past, as is done in abduction [19].

example of the latter would be the following reasoning process: "*Suppose I do not take my umbrella with me. Then I run the risk of getting wet, which I don't want. Therefore I better take my umbrella with me*". The basic inference rule underlying this pattern is modus ponens: if p and p $\rightarrow$ q, then q. However, this example also involves some evaluation of the derived state q: since this state is in conflict with the person's desire, the hypothetical action p (not taking an umbrella) is rejected as being useful.

In the case of criminal behaviour, hypothetical reasoning may play an important role. Various criminologists describe criminals as rational agents that perform careful deliberation of alternatives before they decide to commit a crime e.g., [8]. During such deliberation, they may derive what is the effect of hypothetical actions on the world as well as on other agents. Concerning the latter, the influence of the peers can be very important. Especially adolescents are often put up to a certain action due to peer pressure. For them, it may be very important to belong to a group and have the respect of that group. A theory that is relevant while studying this phenomenon is the theory of *social cognition* by Mead [22]. According to Mead, social cognition arises in situations where people deliberate about planned actions that have impact on other people. Social cognition is a process that is similar to reciprocal role-taking between interactants, except that it occurs in the mind between two phases of the self, namely the *objective* and *acting* self. The objective self consists of a line of actions. This part of the self derives the possible consequences of the planned action and forms a view of the self performing this action from the standpoint of others. Next, the acting self reacts to this: it either responds positively to the objective self and allows the action to be performed or responds negatively and blocks it [21, 22].

In this paper, the ideas of hypothetical reasoning and social cognition are formalised, thereby creating a mechanism that evaluates the different aspects of a certain hypothetical action. These aspects may involve factors that have a biological background (e.g., the desire for aggressiveness) as well as factors that have a more rational nature (e.g., the desire to be appreciated by friends). For example: "*Suppose I attack that person. Then this will satisfy my desire for aggressiveness and my desire to be appreciated by my friends. Therefore I decide to attack that person*".

In the next section, this mechanism will be integrated with [3]'s model for decision making. This model is based on the *rational choice* theory within criminology; e.g. [8], which describes crime as the result of a deliberation process in which pros and cons are weighed. For example, an offender may decide to risk breaking the law, after assessing the chances of getting caught, the expected penalty, the value to be gained by committing the act, his or her immediate need for that value and the appreciation by his or her friends.

3 Simulation Model

In order to develop the model for criminal decision making, the modelling language LEADSTO [5] has been used. This language integrates qualitative, logical aspects and quantitative, numerical aspects, which allows the modeller to exploit both logical and numerical methods for analysis and simulation. Since the domain under consideration involves both qualitative aspects (e.g. decisions to perform a certain action) and quantitative aspects (e.g., utilities of a certain action), it is an appropriate modelling choice. However, in this paper LEADSTO is mainly used as a modelling vehicle. It is not claimed that it is the only possible approach[2].

In LEADSTO, direct temporal dependencies between two state properties in successive states are modelled by *executable dynamic properties.* The LEADSTO 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 the leads to language 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(a, d), expressing that agent a has desire d, or a quantitative, numerical format such as an expression has_value(x, v), expressing that variable x has value v.

As mentioned in the introduction, the basis of the decision making model is a BDI-model, cf. [12, 28]. 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 fulfil this desire. The BDI-model was specified by LEADSTO rule R1 and R2:

R1 Desire d combined with the belief that a certain action ac will lead to the fulfillment of that desire will lead to the intention to perform that action.
$\forall$a:AGENT $\forall$d:DESIRE $\forall$ac:ACTION
desire(a, d) $\wedge$ belief(a, satisfies(ac, d)) $\twoheadrightarrow_{0, 0, 1, 1}$ intention(a, ac)

In the presented model for criminal behaviour, desires may be complex states, which are composed of multiple sub-desires. These sub-desires have been determined based on interviews with domain experts. In the case study presented below, a desire of a criminal is composed of sub-desires for the following aspects: high gain, low loss, negative feelings, actions with strong stimuli, aggressiveness, and appreciation by friends (see [3] for details of the selected aspects). Such a composed desire is represented below as d(hg, ll, ne, ass, das, af). In addition, the following rule was used:

R2 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:AGENT $\forall$ac:ACTION
belief(a, opportunity_for(ac)) $\wedge$ intention(a, ac) $\twoheadrightarrow_{0, 0, 1, 1}$ performed(a, ac)

However, to assess and compare different options, and select a best option, as an extension to this basic BDI-model, personal *utilities* are to be assigned and combined, addressing the degree to which an action satisfies a desire of an individual. Note that these utilities are assumed to be subjective and personal. For example, for a criminal subject, due to his or her specific biological and psychological characteristics, a desire may be quite deviant from what is commonly considered as the rational norm. For this

[2] For example, other modelling languages that have similarities with LEADSTO (and its super-language TTL) are Situation Calculus [29], and Event Calculus [18]. For a comparison with these approaches, see [4].

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, introduce an aspect utility v_i for any possible action ac by

has_aspect_utility(ac, x_1, v_1)

…

has_aspect_utility(ac, x_k, v_k)

where v_i expresses the expected utility for each aspect x_i of the considered action ac, normalised between 0 (minimal utility) and 1 (maximal utility). For example,

has_aspect_utility(fight, desire_for_aggressiveness_satisfied, 0.9)

indicates that the action of fighting contributes much to satisfaction of the desire 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**:** $f(v_1, \ldots, v_k, w_1, \ldots, w_k) = \sqrt{(w_1 v_1^2 + \ldots + w_k v_k^2)}$
- Manhattan Distance**:** $f(v_1, \ldots, v_k, w_1, \ldots, w_k) = w_1 v_1 + \ldots + w_k v_k$

The LEADSTO property for combination is:

R3 ∀a:AGENT ∀ac:ACTION ∀x_1,...,x_k:ASPECT ∀v_1,...v_k,w_1,...w_k:REAL
belief(a, has_aspect_utility(ac, x_1, v_1)) ∧ … ∧ belief(a, has_aspect_utility(ac, x_k, v_k)) ∧
weight_factor(x_1, w_1) ∧ … ∧ weight_factor(x_k, w_k) $\twoheadrightarrow_{0,\,0,\,1,\,1}$
belief(a, has_utility(ac, d(x_1,..., x_k)), f(v_1, …, v_k, w_1, …, w_k)))

Next, the choice process is formalised. This is done in two steps. First, R1 is replaced by R1a, R1b, and R1c:

R1a Desire d combined with the belief that a certain action ac will lead to the fulfillment of d with utility u (≥c) will lead to the consideration of a as a possible intention option.
∀a:AGENT ∀d:DESIRE ∀ac:ACTION ∀u:REAL
desire(a, d) ∧ belief(a, has_utility(ac, d, u) ∧ u≥c) $\twoheadrightarrow_{0.2,\,0.2,\,1,\,1}$ is_intention_option(a, ac, 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:

R1b If ac1 and ac2 are both intention options, but ac2 has a higher utility, then ac1 is ruled out by the agent as an intention option.
∀a:AGENT ∀ac1,ac2:ACTION ∀u1,u2:REAL
is_intention_option(a, ac1,u1) ∧ is_intention_option(a, ac2,u2) ∧ u1<u2 $\twoheadrightarrow_{0,\,0,\,1,\,1}$
ruled_out_intention_option(a, ac1, u1)

R1c Eventually, an intention option that is not ruled out is selected as final intention.
∀a:AGENT ∀ac:ACTION ∀u:REAL is_intention_option(a, ac, u) ∧
not ruled_out_intention_option(a, ac, u)
$\twoheadrightarrow_{0,\,0,\,1,\,1}$ intention(a, ac)

The model presented so far enables an agent to deliberate between possible alternative actions ac, based on its desire d. However, it assumes that the knowledge about the extent to which actions satisfy desires (i.e., state properties of the form belief(a, has_utility(ac, d, u))) is given. Therefore, the following mechanism is proposed to derive such relations dynamically. The reasoning starts by assuming that the agent performs a certain (hypothetical) action ac (assumed(a, ac)). Next, knowledge about dependencies between world states (state properties of the form belief(a, leads_to(ac, y, v))) is exploited to predict the certainty with which actions lead to world states:

R4 If an agent currently assumes ac, and believes that ac leads to x with certainty v, then it will predict that ac leads to x with certainty v.
∀a:AGENT ∀ac:ACTION ∀x:INFO_ELEMENT ∀v:REAL
assumed(a, ac) ∧ belief(a, leads_to(ac, x, v)) $\twoheadrightarrow_{0,\,0,\,1,\,1}$ predicted(a, leads_to(ac, x, v))

R5 If an agent currently predicts that ac leads to x with certainty v, and believes that x leads to y with certainty w, then it will predict that ac leads to y with certainty v*w.
∀a:AGENT ∀ac:ACTION ∀x,y:INFO_ELEMENT ∀v,w:REAL
predicted(a, leads_to(ac, x, v)) ∧ belief(a, leads_to(x, y, w))
$\twoheadrightarrow_{0,\,0,\,1,\,1}$ predicted(a, leads_to(ac, y, v*w))

Finally, in case a prediction is made of an aspect that is relevant for the agent (because it has a desire for that aspect), then the probability of that aspect is taken as aspect utility.

R6 Predictions about relevant aspects are used as aspect utilities.
∀a:AGENT ∀ac:ACTION ∀x:INFO_ELEMENT ∀v:REAL
predicted(a, leads_to(ac, x, v)) $\twoheadrightarrow_{0,\,0,\,1,\,1}$ belief(a, has_aspect_utility(ac, x, v))

The complete utility-based decision model is depicted graphically in Figure 1 (where the variables denoting the reasoning agent a have been left out, for simplicity). The circles denote state properties, and the arrows denote dynamic (LEADSTO) properties. Notice that the state properties of the type desire(...) can be generated by detailed submodels with domain-specific knowledge about particular types of criminals (e.g., as presented in [2]), which are outside the scope of this paper.

4 Example Simulation Trace

Based on the model shown above, a number of simulation experiments have been performed to test (for a simple scenario) whether it shows the expected behaviour. In this section, an example simulation trace is described in detail.

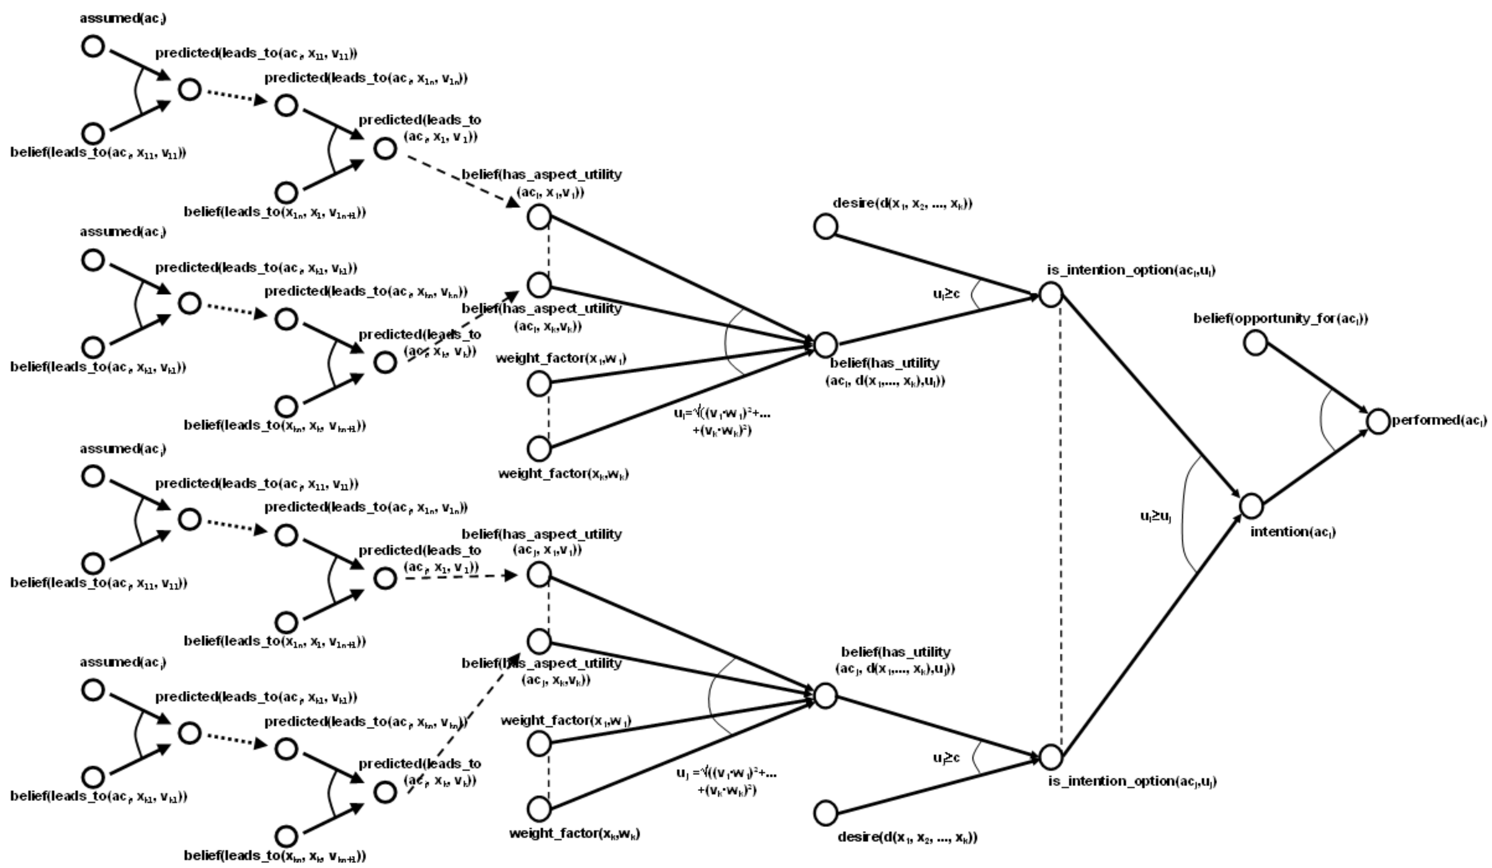

Figure 1: Overview of the Decision Making Model

The example scenario involves a boy A, who hangs out with his friends at school. These friends have the reputation of being bad boys. They like to harass their fellow students and make trouble. The friends have a negative influence on A, who actually is quite a nice boy. To be accepted a bit more by his friends he wants to impress them. An opportunity occurs; one of their fellow students arrives on his bike. Boy A thinks that attacking the student and taking his bike might have a good impact on the appreciation by his friends. Therefore he starts weighing the pros and cons. Boy A's individual weight factors for the different aspects involved are as follows: {high_gain = 0.14, low_loss = 0.17, negative_feelings = 0.08, actions_with_strong_stimuli = 0.19, desire_for_aggressiveness_satisfied = 0.17, appreciation_by_friends = 0.25}.

These weight factors, which add up to one, show the relative importance of each aspect for A. For example, the weight factor of high gain is 0.14, i.e., A does not find high gain very important. Appreciation, however, has a weight factor of 0.25: A finds it really important that his friends appreciate his behaviour. A is deliberating whether or not to perform the action (so to attack the student or do nothing).

Table 1 presents A's beliefs about dependencies between different world states. For each dependency, the first column shows the antecedent, the second column shows the consequent, and the third column shows the certainty A has about the dependency. For example, the fourth row is formalised via the state property:

belief(criminal,
    leads_to(belief(friend, performed(criminal, attack_student)), appreciation_by_friend, 0.9))

which expresses that the criminal believes that, when his friend knows that he attacked the student, he will probably appreciate this (with a probability of 90%).

Table 1: Agent A's Beliefs about World State Dependencies

antecedent	consequent	certainty
performed(criminal, attack_student)	world_state(performed(criminal, attack_student)	0.95
world_state(performed(criminal, attack_student)	observes(friend, performed(criminal, attack_student)	0.9
observes(friend,performed(criminal, attack_student)	belief(friend, performed(criminal, attack_student)	1.0
belief(friend,performed(criminal, attack_student)	appreciation_by_friend	0.9
world_state(performed(criminal, attack_student)	world_state(stolen(criminal, bike))	0.9
world_state(performed(criminal, attack_student)	world_state(performed(criminal, not_get_caught))	0.7
world_state(performed(criminal, attack_student)	experiences(student, suffering)	0.9
world_state(performed(criminal, attack_student)	experiences(criminal, adrenalin_rush)	1.0
world_state(performed(criminal, attack_student)	world_state(performed(criminal, hit_student)	0.8
performed(criminal, nothing)	world_state(performed(criminal, nothing)	0.95
world_state(performed(criminal, hit_student)	desire_for_aggresiveness_satisfied	0.9
world_state(stolen(criminal, bike))	high_gain	0.3
world_state(performed(criminal, not_get_caught))	low_loss	0.3
experiences(student, suffering)	negative_emotions	0.8
experiences(criminal, adrenalin_rush)	actions_with_strong_stimuli	0.6

Figure 2 shows (part of) an example simulation trace that was generated on the basis of these settings. The example trace is the output of the LEADSTO simulation environment mentioned above [5]. Here, time is on the horizontal axis and state properties are on the vertical axis. A dark line indicates that a state property is true.

As can be seen in Figure 2, agent A (represented as criminal in this simulation trace) has the opportunity to attack the student (belief(opportunity_for(performed(criminal, attack_student)))) or to do nothing (belief(opportunity_for(performed(criminal, do_nothing)))) at the beginning of the simulation. At time point 5, the criminal starts his deliberation process by assuming that he attacks the student. Then, he derives that this will lead to a probability of 85.5% that his friends will see this action (time point 7). He also derives that the attack of the student will probably lead to a stolen bike (85.5%), not getting caught (66.5%), suffering by the student (85.5%), an adrenalin rush (95%) and the student getting hit (76%). One step further, this might (but is not very certain to) lead to high gain (25.6%) and low loss (19.9%). It will probably lead to negative emotions (68%), the fulfillment of an action with strong stimuli (57%) and a satisfied desire for aggressiveness (68%), see time point 8. Furthermore, he derives that his friends will appreciate the action with 77% probability. The combination of these aspects makes that the student concludes that the attack is a possible action (is_intention_option) with a utility of 0.54. After that, A starts to consider the option of doing nothing. This decision making process is similar to the process of calculating the possible action of attacking the student (and was therefore left out of Figure 2). Based on this process, the student concludes that doing nothing has a utility of 0.10. Since the action of attacking the student has a higher utility, the criminal decides to attack the student (at time point 22) and actually attacks the student (at time point 23).

Figure 2: Example Simulation Trace

All in all, this and similar scenarios have been simulated various times, using different parameters for the weight factors and the certainty of the dependencies involved. The computational time to generate such simulation traces was not more than a couple of seconds per trace. Due to space limitations, the details of these simulations are not shown here. Nevertheless, this case study indicated that the model enables agents to deliberate about the consequences of their actions, thereby weighing their pros and cons in the context of their (biologically and psychologically determined) desires, and eventually making a rational choice. It also shows that agents may make different decisions, based on personal characteristics and priorities.

The resulting simulation traces have been shown to expert criminologists, who confirmed the plausibility of the produced patterns, and their correspondence to literature such as [8, 23]. Although this is not an exhaustive validation, it is an indication that the model produces expected behaviour.

5 Discussion

This paper presented an agent-based model for decision making, which integrates personal biological and psychological aspects (e.g. high levels of testosterone) with rational utility-based reasoning. The model takes a BDI-based approach, where generation of desires is based on the personal characteristics, and generation of intentions is based on the rational reasoning. Compared to an earlier version [3], the current paper specifically adds a mechanism to derive knowledge that connects certain actions to desires. This is done by means of hypothetical reasoning: the agent imagines performing an action and applies a forward reasoning strategy in order to derive the consequences, and to what extent these fulfil the agent’s desires.

The model has been implemented in the LEADSTO environment [5]. For a case study in criminal behaviour, various simulation experiments have been performed. These simulations have pointed out that the model enables agents to reason about the consequences of their actions, and to take this into account during decision making.

The presented model was inspired by related approaches in the area of Artificial Intelligence. In recent years, several authors have proposed agent-based models for reasoning and decision making (which often also combine rational and non-rational aspects), such as in [14, 31]. 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 arousal). Furthermore, the presented model has some similarities with standard BDI-based agent modelling frameworks, such as AgentSpeak [27] or Jadex [24]. In further research, it may be investigated whether the proposed decision making model can be implemented in such frameworks.

Concerning further future work, these are several possibilities. As mentioned in the introduction, the presented model can be useful from two perspectives. First, it may be used to get more insight in the process of criminal decision making itself: why do certain people commit certain crimes? To answer such questions, the model has to be fed with more detailed domain knowledge from criminological and psychological literature, and specific experiments have to be designed to study typical questions described in the literature. Some initial steps in this direction have been made, see also [2, 3]. Second, the model can be used to develop *virtual agents* cf. [25] that are able to make decisions based on hypothetical reasoning, and more specifically, to develop ‘criminal agents’. For example, such criminal agents may play a role in (serious) games and computer-generated virtual stories e.g., [7]. Since the LEADSTO language has an intuitive (causal relationship-based) format, and is independent of a particular implementation language, this step is relatively straightforward, see also [6]. When doing that, also LEADSTO’s possibilities to specify non-deterministic reasoning rules can be exploited. Work is currently in progress to elaborate these ideas in detail.

References

1. Baron, J. (1988). *Thinking and Deciding.* Cambridge University Press, 1988.
2. Bosse, T., Gerritsen, C., and Treur, J. (2007). Integration of Biological, Psychological and Social Aspects in Agent-Based Simulation of a Violent Psychopath. In: Shi, Y., Albada, G.D. van, Dongarra, J., and Sloot, P.M.A. (eds.), *Proceedings of the Seventh International Conference on Computational Science, ICCS'07*. Lecture Notes in Computer Science, vol. 4488, Springer Verlag, 2007, pp. 888-895.
3. Bosse, T., Gerritsen, C., and Treur, J. (2007). Integrating Rational Choice and Subjective Biological and Psychological Factors in Criminal Behaviour Models. In: Lewis, R.L., Polk, T.A., and Laird, J.E. (eds.), *Proceedings of the 8th International Conference on Cognitive Modeling, ICCM'07*. Taylor and Francis, 2007, pp. 181-186.
4. Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J. (2009). Specification and Verification of Dynamics in Agent Models. *Int. Journal of Cooperative Information Systems*.
5. Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2007). A Language and Environment for Analysis of Dynamics by SimulaTiOn. *Int. Journal of AI Tools*, vol. 16, issue 3, pp. 435-464.
6. Bosse, T., and Lange, F.P.J. de (2008). Development of Virtual Agents with a Theory of Emotion Regulation. In: *Proc. of the Eighth IEEE/WIC/ACM International Conference on Intelligent Agent Technology, IAT'08*. IEEE Computer Society Press, 2008, in press.
7. Cavazza, M., Charles, F., Mead, S. (2003). Interacting with virtual characters in interactive storytelling. In: Alonso, E., Kudenko, D., and Kazakov, D. (eds.), Adaptive Agents and Multi-Agent Systems. LNAI, volume 2636, Springer Verlag, 2003, pp. 318–325.
8. Cornish, D.B., and Clarke, R.V. (1986). *The Reasoning Criminal: Rational Choice Perspectives on Offending*. Springer Verlag.
9. Damasio, A.R. (1994). *Descartes' error: Emotion, reason, and the human brain*. NY: Putnam.
10. Evans, J.S.B.T. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences* 7, pp. 454-459.
11. Fishburn, P. C. (1970). *Utility Theory for Decision Making*. NY: John Wiley and Sons.
12. Georgeff, M. P., and Lansky, A. L. (1987). Reactive Reasoning and Planning. In: *Proc. of the Sixth National Conf. on Artificial Intelligence*, AAAI'87. Menlo Park, California. American Association for Artificial Intelligence, 1987, pp. 677-682.
13. Hanks, S. and McDermott, D. (1987). Nonmonotonic logics and temporal projection. *Artificial Intelligence*, 33, pp. 379-412, 1987.
14. Jiang, H., Vidal, J.M., and Huhns, M.N. (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, pp. 38-40.
15. Jonker, C.M., and Treur, J. (2003). Modelling the Dynamics of Reasoning Processes: Reasoning by Assumption. *Cognitive Systems Research,* vol. 4, 2003, pp. 119-136.
16. Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, pp. 263-291, 1979.
17. Keeney, R., and H. Raiffa. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, 1976.
18. Kowalski, R., and Sergot, M. (1986). A logic-based calculus of events, New Generation Computing, 4, pp. 67-95.

19. Lifschitz, V., (1997). On the logic of causal explanation. In: Artificial Intelligence, volume 96, issue 2, pp. 451-465.
20. Loewenstein, G. F., Weber, E. U., Hsee, C. K., and Welch, E, S. (2001). Risk as feelings. *Psychological Bulletin*, vol. 127, pp. 267-286.
21. Matsueda, R.L., and Kreager, D. (2004). *Social Cognition, Identity, and Deterrence: a model of criminal decisions*. In: Proc. of the Annual Meeting of the Am. Sociological Association, 2004.
22. Mead, G.H. (1934). *Mind, Self and Society: from the standpoint of a social behaviorist*. Edited by Charles W. Morris. Chicago: University of Chicago, 1934.
23. Moir, A., and Jessel, D. (1995). *A Mind to Crime*: *the controversial link between the mind and criminal behaviour.* London: Michael Joseph Ltd; Penguin.
24. Pokahr, A., Braubach, L., and Lamersdorf, W. (2005). Jadex: A BDI reasoning engine. In: Bordini, R.H., Dastani, M., Dix, J., and El Fallah Seghrouchni, A. (eds). *Multi-Agent Programming: Languages, Platforms and Applications*. Springer Verlag, 2005, pp. 149–174.
25. Prendinger, H., Lester, J., and Ishizuka, M. (eds.) (2008). *Intelligent Virtual Agents. Proceedings of the 8th International Conference on Intelligent Virtual Agents, IVA'08*. Springer LNAI, vol. 5208.
26. Raiffa, H. (1982). *The Art and Science of Negotiation*. Harvard University Press, Cambridge, MA, 1982.
27. Rao, A. S. (1996). AgentSpeak(L): BDI agents speak out in a logical computable language. In: de Velde, W.V. and Perrame, J., (eds.), *Agents Breaking Away: Proceedings of the Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW'96),* pages 42–55. Springer Verlag. LNAI, Volume 1038.
28. Rao, A.S. and Georgeff, M.P. (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, pp. 473-484.
29. Reiter, R. (2001). *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*, Cambridge MA: MIT Press.
30. Rescher, N. (1964). *Hypothetical reasoning*. Amsterdam: North-Holland.
31. Silverman, B.G. (2001). More Realistic Human Behavior Models for Agents in Virtual Worlds: Emotion, Stress, and Value Ontologies. T.R., Philadelphia, PA, University of Penn./ACASA.