Modelling a Fighter Pilot's Intuition in Decision Making on the Basis of Damasio's Somatic Marker Hypothesis

Mark Hoogendoorn[1], Robbert-Jan Merk[2], Jan Joris Roessingh[2], Jan Treur[1]

[1]Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081, 1081 HV Amsterdam, The Netherlands
{mhoogen, treur}@cs.vu.nl
http://www.few.vu.nl/~{mhoogen, treur}
[2]NLR , Training, Simulation & Operator Performance
Anthony Fokkerweg 2, 1059 CM Amsterdam, The Netherlands
{merkrj, roess}@nlr.nl

Damasio's Somatic Marker Hypothesis" (Damasio, 1994) provides a neuro-cognitive account on how humans manage to solve complex decision problems in a relatively efficient way. Damasio proposes that specific centers in the brain, which relate to feeling certain body states resulting from emotional responses on a situation, generate so-called 'Somatic Markers' that help to make a selection of viable options for action. This process involuntarily guides the experts' decision and provides him with the sense that a complex decision was taken care of by gut feeling, or rather intuition.

Fighter pilots very much depend on the type of aggregated experiences that constitute their intuition. It enables them to stay separated from their enemies and attack them from the right angle, at the same time coordinating with their wingmen in the formation, and all of this in a few seconds. Their considerations of the tactical situation, the possible maneuvers and actions and the desired outcome of the situation do not proceed rationally. These decisions and evaluations are generated subconsciously and intuitively.

Fighter pilots obtain this competency gradually through training, in a 'live' or a simulated environment. In the latter type of environment, the opponent aircraft and their operators can be represented by Computer Generated Forces (CGFs). In so far these CGFs adapt themselves to the tactical environment, they can be considered virtual intelligent agents. The idea is that such agents behave human-like. They observe the environment, orientate themselves in this environment, decide on what to do next and perform actions in a way that humans would do, including opportunistic, erroneous, untraceable and sometimes surprisingly brilliant behavior that is so characteristically human.

To discover the essence of realistic role behavior, expert knowledge of fighter pilots is used. In addition, available knowledge from Cognitive Science is translated into specific models that underlie these virtual agents. Such models may include elements of intuitiveness, creativity, opportunism, autonomy, and situation awareness, but also elements of exhaustion, task saturation, and psychological and physical stress.

Present models of decision making are predominantly based on the assumption that the basis for decisions is merely rational. Kahneman and Tversky (1976) already pointed out that even for fairly simple decisions humans do not operate on a purely rational basis. Hence, a model on the basis of Damasio's Somatic Marker Hypothesis has been designed to account for the role of intuitiveness, emotions and feelings in fighter pilots' decision making.

This computational model has been modelled in the declarative hybrid modelling language LEADSTO which integrates temporal logical and numerical capabilities. The model focuses on key aspects of the Somatic Marker Hypothesis. For example, a central feature of the model is Damasio's claim that somatic marking helps to reduce the number of options in order to make a cost-benefit analysis more manageable.

The validity of the model is demonstrated by means of simple tactical situations taken from the air-to-air fighter domain that are used as scenarios for the LEADSTO software environment in which simulations with the model have been performed.

INTRODUCTION

Domain

This research is positioned in the flight simulation domain. More specifically, mission tactical simulations are used in the training and mission rehearsal of fighter pilots to acquire competency in dealing with adversaries in tactical situations. The adversary, for example an enemy fighter, is implemented in the simulation as a Computer Generated Force (CGF). A CGF is a simulated entity that interacts with other entities (either real and/or simulated). The behaviour of the real-life version of that entity, e.g. the enemy fighter aircraft, is mimicked by a computer model, which is an integral part of the simulation. It is obvious that the training value of the tactical exercise will to a large extent depend on the (required) realism of the CGF. CGFs that represent weapon platforms controlled by human operators must therefore include a model of these human operators. Such a model should provide the interactive behavior that can be expected and observed from real-life human operators in a convincing way. This means that the CGF needs to act as an 'intelligent agent', i.e. as if it understands the consequences of its actions, just like a competent fighter pilot would understand.

Three main contributory factors to mission success are considered as fundamental for the behaviour of the intelligent agent, i.c. a fighter pilot executing his mission. These factors are (1) lethality, i.e. the ability of the weapon platform to kill the opponent in a given situation, (2) survivability, i.e. the ability to remain mission capable after the engagement with an opponent, and (3) resource control, i.e. the efficient and effective use of resources, such as weapons, fuel, sensor intelligence and C2 capabilities. The terms 'efficient' and 'effective' refer to the planned mission goals. Surviving the mission and return home safely is assumed to be a leading goal in any mission. For example, in an air-to-ground mission an F-16 pilot might drop-off its bomb load (resource control) to cope with the sudden appearance of an enemy formation. This decision effectively increases his survivability, but decreases his chances to fulfill his original mission objectives to zero. A basic assumption of the proposed model is that the pilot makes an assessment of these three factors, intuitively or explicitly, in engagement with enemy forces.

Goal of the research

The CGF/ agent with which the trainee pilot engages must thus have certain intelligence to turn the mission tactical simulation into an effective training exercise. Currently, CGFs in the fighter pilot domain are restricted in this respect: CGF behavior is scripted rather than situation oriented or the CGF provides sensible behavior in a limited tactical situation only. One reason for the lack of realistic behaviour is that populating a rule-based reasoning engine with data requires a lot of effort. This research tests an approach which is based on 'somatic marking' to make a down-select of the viable behavioral options. Somatic marking is based on Somatic Markers, i.e. intuitive feelings, such as fear, aggressiveness, anger, anticipation and surprise. In the resultant agent, intuition becomes an integral part of its decision making, rather than an emotional layer around a core of rational decision making, as can be found in some agents (e.g. CoJACK, Bittner, Busetta, Evertsz and Ritter, 2008). The agent can be trained in an off-line simulation, using valuation of decision alternatives on the basis of both Somatic Markers and rational arguments. This results in generalizable behaviour on the basis of fewer and less complex rules which the agent can use to its advantage when plugged into human-in-the-loop simulation. The current research aims at prototyping an example of such agent.

Scope

In this first approach, the focus is on cognitive aspects, particularly tactical decision making, for the current purposes leaving aside perceptual-motor processes. Before going into detail, a simple case from a tactical mission is discussed, denoted 'surface-to-air threat handling'.

CASE

To demonstrate the principles that underlie intuitive decision making, a simplified case of an agent is presented, in the role of a fighter pilot, who will be taught how to avoid being shot down by ground-based anti-air weapons.

Different kinds of these anti-air weapons are considered:

1. Radar-guided Surface-to-Air Missiles (SAMs)
2. Infra-Red(IR)-guided SAMs
3. MAN-Portable Air-Defense (MANPAD) missiles, based on Line-of-Sight (LoS) guidance
4. Anti-Air Artillery (AAA)

For each of the weapon systems the agent must be taught to execute different tactics in order to avoid being shot down. Note that the described tactics are a gross simplification of the actual tactics employed in reality.

A site equipped with radar-guided SAMs uses radar to track the aircraft and guide the SAM towards it. Once the agent becomes aware that he is in range of the SAM, the tactic is to perform a "beam" manoeuvre, that is, to fly perpendicular to the radar beam as to minimize radar visibility. Also, the agent can defend himself to dispense chaff (a cloud of thin pieces of aluminium) to distract radar-guided missiles.

IR-guided SAMs are guided by the heat, especially the engine exhaust, emitted by the aircraft. The tactic is to dispense heat-emitting flares that confuse the heat-seeking SAMs.

With MANPAD missiles, a human operator guides the missile to the target using a laser beam or radio control. The range of MANPAD missiles is short and if LoS is lost the missile cannot find the target on its own. Once detected being under fire by a MANPAD (through a missile approach warning in the cockpit), the tactic is to break LoS through an evasive manoeuvre with the appropriate parameters.

The idea behind the use of AAA is that if enough rounds of ammunition are shot into the air, some are bound to hit the aircraft. This makes AAA quite lethal for aircraft flying at low attitude. Once the agent is aware of enemy AAA, the tactic is to stay out of range / increase altitude.

Choice of tactics, outcome and valuation

The agent has several options in terms of tactics. However, there are only two possible outcomes with each choice, either the aircraft is hit or missed. The outcome thus depends on the tactics chosen in combination with the type of surface-to-air threat. Table 1 defines the relationships between the chosen tactical option, anti-air system and outcome.

It is assumed that unless the agent executes a proper counter tactic, it will be hit. Therefore, ignoring a threat always results in a hit. Likewise, the agent can choose to re-plan or abort the mission such that he doesn't enter the ranges of anti-air threats. In this case, the other four tactical options are suited against one particular type of threat only.

Table 1: Anti-air threat / tactics combinations and their outcomes

Tactical option		Type of anti-air threat			
		Radar SAM	IR SAM	MAN-PAD	AAA
1	Ignore and continue	Hit	Hit	Hit	Hit
2	Beam and chaff	Miss	Hit	Hit	Hit
3	Flares	Hit	Miss	Hit	Hit
4	Increase altitude	Hit	Hit	Miss	Miss
5	Break LoS	Hit	Hit	Miss	Hit
6	Re-plan/ abort	Miss	Miss	Miss	Miss

Table 2: Attribution of rational utility to the tactical options

Tactical Option		Rational Utility
1	Ignore and continue	1.0
2	Beam and chaff	0.5
3	Flares	0.5
4	Increase altitude	0.7
5	Break LOS	0.8
6	Re-plan/ abort	0.3

To teach the agent in making the tactical decisions, a rational utility value is attributed to each tactical option (see Table 2).

Utility of each tactical option is based upon two factors: (1) the amount of mission resources (such as: time, fuel, chaff and flares) involved (the more resources are consumed, the lower the utility), (2) how it changes the chance to successfully complete the mission (the lower the chance, the lower the utility). Obviously, the option “ignore and continue” does not use additional resources and is the “best” way to continue the mission, if survivability is not an issue. In contrast, the costs of re-planning or aborting the mission are high, possibly discarding the mission goals.

Survivability considerations are not taken into account in the attribution of rational utility. In this case it is assumed that survivability is taken into account intuitively, that is, through Somatic Markers, which is further explained in the next sections.

THEORY

Decision making and experience

Decision making usually involves expectations about possible consequences of decision options and uncertainty about them. Traditionally the literature on the analysis of decision making was dominated by the Expected Utility Theory; see, for example, von Neumann and Morgenstern, 1944; Friedman and Savage, 1948; Arrow, 1971; Keeney and Raiffa, 1976. Here the process of decision making takes place by calculating the expected utilities for all of the options and choosing the option with highest expected utility. The expected utilities themselves are determined based on the probabilities of the possible outcomes for the option when chosen, and the gain or loss for that outcome, thus founding the approach in probability theory. This approach to decision making can be considered to aim for an idealised rational approach, where, for example, emotions or biases play no role. As a model for practical human decision making the Expected Utility Theory has been strongly criticized, as humans are usually bad in estimating probabilities, and also may allow emotions and biases to play a role in a decision process, as is found in several experiments; see, for example (Tversky and.Kahneman, 1974; Kahneman and Tversky, 1979).

Contrasting with the aim of the Expected Utility Theory to ban emotions from decision making, Damasio (1994) observed surprisingly bad decision making behaviour in patients with damage of brain regions related to body mapping and regulation and feeling emotions (patients with certain kinds of prefrontal damage and with compromised emotions). They often keep on considering different options without choosing for one of them. This has led Damasio to the view that decision making inherently depends on emotions felt, which according to his perspective relate to sensed body states or preparations for body states (Damasio, 1994). Part of this process concerns a nonconscious biasing mechanism pointing to the better decision options. His theory is known as the *Somatic Marker Hypothesis*. Damasio explains the name of his theory as follows:

> ‘Because the feeling is about the body, I gave the phenomenon the technical term *somatic* state (“soma” is Greek for body); and because it “marks” an image, I called it a *marker*. Note again that I use *somatic* in the most general sense (that which pertains to the body) and I include both visceral and nonvisceral sensation when I refer to somatic markers.’ (Damasio, 1994, p. 173)

This theory consists of two main ideas: (1) the way in which Somatic Markers affect decisions, and (2) the way in which Somatic Markers depend on past experiences.

(1) The way in which Somatic Markers affect decisions

If a decision must be made between options which can lead to potentially harmful or advantageous outcomes, each of such an option induces a somatic response which is sensed as a feeling and used to mark the option outcome, thus signalling its danger or advantage. For example, when a negative Somatic Marker is linked to a particular option outcome it serves as an alarm signal for that particular option. Similarly, a positive Somatic Marker serves as an encouragement to choose for that option. Damasio describes the consequence of a Somatic Marker in the following way:

> ‘the somatic marker (..) forces attention on the negative outcome to which a given action may lead, and functions as an automated alarm signal which says: Beware of danger ahead if you choose the option which leads to this outcome. The signal may lead you to reject, *immediately*, the negative course of action and thus make you choose among other alternatives. The automated signal protects you against future losses, without further ado, and then allows you *to choose from among fewer alternatives*’ (Damasio, 1994, p. 173)

‘In short, *somatic markers are a special instance of feelings generated from secondary emotions*. Those emotions and feelings *have been connected by learning to predicted future outcomes of certain scenarios*. When a negative somatic marker is juxtaposed to a particular future outcome the combination functions as an alarm bell. When a positive somatic marker is juxtaposed instead, it becomes a beacon of incentive. This is the essence of the somatic marker hypothesis. (..) on occasion somatic markers may operate covertly (without coming to consciousness) and may utilize an ‘as-if-loop’. (Damasio, 1994, p. 174)

(2) The way in which Somatic Markers depend on past experiences

The way in which Somatic Markers are associated to decision options in a given situation depends on previous experiences with options chosen in similar circumstances. For example, the pain or joy experienced as a consequence of the outcome for a certain option that was chosen in the past has been stored in memory and automatically pop up (are felt again) when similar circumstances and options may occur. How Somatic Markers relate to past experiences is described as follows:

‘Somatic markers are thus acquired through experience, under the control of an internal preference system and under the influence of an external set of circumstances which include not only entities and events with which the organism must interact, but also social conventions and ethical rules. (Damasio, 1994, p. 179)

This element of Damasio’s theory shows how based on experience ‘intuition’ or ‘gut feeling’ is created which aids the decision process in an automatic manner. This makes the theory useful for decision processes where such aspects play an important role, which is the case for the domain of pilot behaviour considered here.

MODEL DESCRIPTION

Modelling assumptions

To formalise Damasio’s Somatic Marker Hypothesis an approach was chosen based on the following assumptions.
General assumptions:

- For any given type of emotion, (persistent) Somatic Markers are related to combinations of (situational) contexts and decision options that are relevant for this context
- At any point in time that a decision has to be made within a given context, somatic evaluation values associated to the options that are relevant for the given context are used
- Somatic Markers and somatic evaluation values are expressed as real numbers between 0 and 1
- Contexts and decision options are expressed as discrete instances

Assumptions on the decision process:

- Within a given context, every decision option gets a somatic evaluation value associated based on the available Somatic Markers for the same context and option
- Decision options with low associated somatic evaluation value are eliminated from further decision processing
- For the remaining decision options after this elimination a (utility-based) rational analysis is made in which the somatic evaluation values serve as biases

Assumptions on adaptation:

- Based on experiences with respect to the outcomes of chosen options for a given context, the Somatic Markers are adapted over time

Modelling approach

The model has been defined as a set of temporal relations between properties of states. The definition of a state as used in this paper is a set of state properties that hold or do not hold at a certain moment. A state property is a conjunction of atoms or negations of atoms.

For example, to describe the state of the weather at monday as sunny, with temperature 25 degrees Celsius and without rain, a state property would be

sunny & not(rain) & temperature(25).

The exact choice for what atoms to use depends on the actual model and domain and is defined by an ontology for that model.

To model dynamics, it is necessary to define the transitions between states. For example, that if it rains at one day, the next day the streets will be wet can be represented formally as a temporal relation

$\text{rain} \twoheadrightarrow_{\text{next day}} \text{state\_of\_streets(wet)}$

This relation applied to the Monday state will make the atom state_of_streets(wet) hold at Tuesday.

In order to obtain an executable formal model, the states and temporal relations between them have been specified in LEADSTO, a temporal language in which the temporal relations can be defined in the form of rules that can be executed. Let α and β be state properties. In LEADSTO specifications 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* h.

For more details of the LEADSTO format, see (Bosse, Jonker, van der Meij & Treur., 2007). As all of the temporal relations used in the model are of the form $\alpha \twoheadrightarrow_{0,0,1,1} \beta$, the notation form $\alpha \twoheadrightarrow \beta$ will be used instead.

The next sections will first provide a high-level overview of the model and then detailed information on the Decision Making and Somatic Marker adaptation parts of the model.

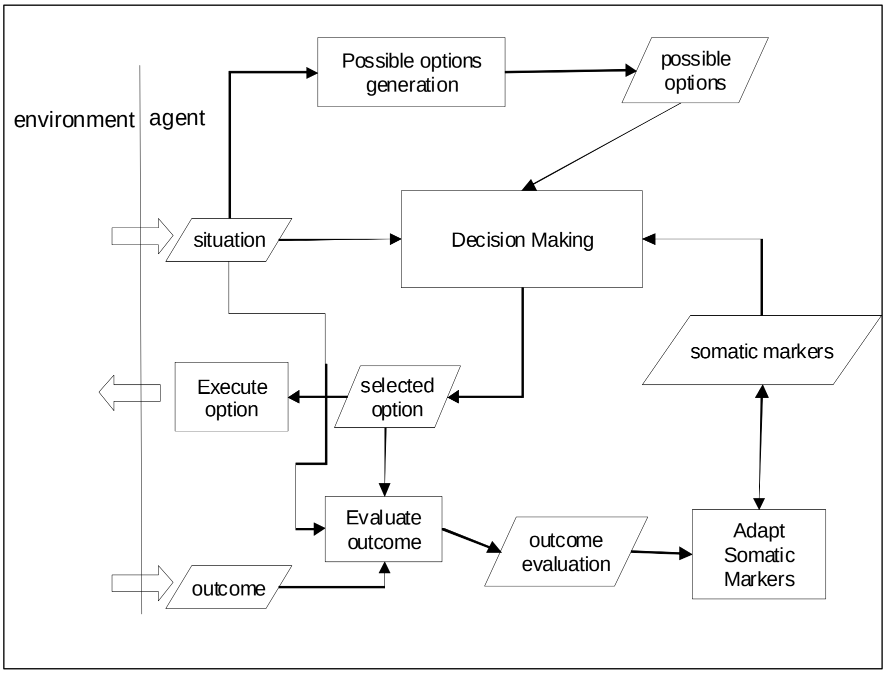

Figure 1: Overview of the model

Overview of the model

The model conceptually consists of a number of processes and transfer of data between these processes. Figure 1 gives an overview of the model, where rectangles are processes and parallelograms are data. This section will describe the model at a high level, the subsequent sections contain more detailed descriptions of the parts of the model.

The central process in the model is the Decision Making process. Its input is the current situation, the list of possible options from which one option is to be selected and the Somatic Markers. For the sake of simplicity, situations and options are represented as single atoms. For more complex cases more complex representations for options and especially situations would be required, but for the case described earlier a simple representation is adequate. The situation is supplied by the environment and can be seen as the result of the agent's perception of its environment. As Damasio's theory does not deal with perception in detail, the model simplifies perception such that the representation of the situation consists of a single atom. For example, in the case described earlier, the agent could encounter an infrared-guided SAM system. In the model the Decision Making process would receive the atom observed(infrared_SAM).

The output of the Decison Process, the selected option, is then executed. Execution of the selected option will result in some change in the environment of the agent and the agent will observe this outcome. This outcome is then evaluated, resulting in a real number between *0* and *1*, where a higher value means a more positive evaluation. The selected option itself is also input for the evaluation process, as the evaluation is about the consequences of this selected option.

The value of the outcome evaluation is then used to adapt the Somatic Markers the agent has. In subsequent decisions the updated Somatic Markers are used.

The Decision Making Process

In the Decision Making process first for each option the Option Evaluation Value is determined. The option with the highest Option Evaluation Value is then selected for execution by a two-step elimination and selection process. Thus the Decision Making process consists of three subprocesses, the Somatic Evaluation Process, Option Elimination and the Rational Analysis. Figure 2 shows the relationships between these processes and the relevant data. The following LEADSTO relations define the dynamics for the input and output of the Decions Making process.

P1 observed(S) ↠
belief_current_situation(S) & decision_making_started

P2 belief_current_situation(S) & not(decision_making_ended) ↠
belief_current_situation(S)

P3 decision_making_ended & selected_option(O) ↠
executed(O)

Property P1 defines the link between perception and the start of the decision making process: the agent stores the observation of the current situation as a belief and start the process. Property P2 is a so-called persistence rule: without it the atom

belief_current_situation(S)

would only hold in the next state. Using P2 the atom holds for all states until the atom decision_making_ended holds. At the end of the Decision Making Process the atom

selected_option(O)

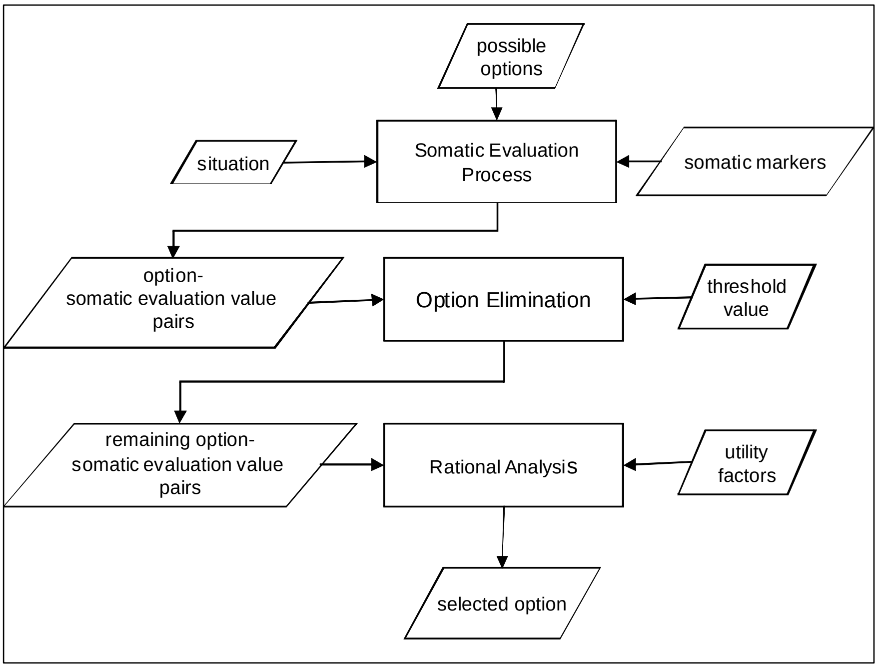

Figure 2: The decision making process in more detail

will hold for some option O. Property P3 defines the fact that a selected option is executed.

P4 belief_current_situation(S) & decision_making_started & possible_option(O) & somatic_marker(S, O, V) ↠ somatic_evaluation(O, V)

P5 somatic_evaluation(O, V) & not(somatic_evaluation_ended) ↠ somatic_evaluation(O, V)

Property P4 defines the way Somatic Evaluation Values are determined. Note that, given the universal variables, this temporal rule is applied for each possible option separately, so that for each option there is a Somatic Evaluation Value. In the model there is currenlty only one type of Somatic Marker, in future work models with multiple types of Soamtic Marker will be explored. Because of this, calculating the Somatic Evaluation Value consists of simply combining an option and the value from the corresponding Somatic Marker. Property P5 is another persistence rule so that the Somatic Evaluation Values are available during the complete Somatic Evaluation process.

P6 somatic_evaluation(O, V) & value(threshold, TH) & V ≥ TH ↠ remaining_somatic_evaluation(O,V) & somatic_evaluation_ended

P7 remaining_somatic_evaluation(O, V) & not(decision_making_ended) ↠ remaining_somatic_evaluation(O, V)

Property P6 defines the Option Elimination process. All the options for which the Somatic Evaluation Value is equal or higher than the threshold are generated, represented by atoms of the form

remaining_somatic_evaluation(O,V)

which form the input for the next step in the Decision Making Process, the Rational Analysis. Property P6 also ends the Somatic Evaluation Process by making the atom

somatic_evaluation_ended

true. This will also end the application of P5 to future states, effectively eliminating all options with a Somatic Evaluation Value below the threshold for further consideration.

The next subprocess is the Rational Analyis. In this process a Rational Utility is calculated for each option for which the atom

remaining_somatic_evaluation(O,V)

holds. The assumption here is that based on the remaining Somatic Evaluation Values and some other input defined here as Utility Factors contribute to determining the utility for each option. In the design of the model there are atoms of the form

option_utility(O,U)

which couple each option O with a real value U between 0 and 1, indicating the utility for that option. More elaborate utility functions are certainly possible but fall outside the scope of this paper.

P8 remaining_somatic_evaluation(O, V) & option_utility(O,U) ↠ option_evaluation(O, (U+V) / 2)

P9 option_evaluation(O, V) & not(decision_making_ended) ↠ option_evaluation(O, V)

In P8 for each all remaining option the Option Evaluation is finally determined. This value is taken as a weighted average between the Somatic Evaluation Value and the Option Utility (for the sake of simplicity for the moment the weights are taken 0.5). The Option Evaluation also persists until the end of the Decision Making Process.

At the end of the Rational Analysis process, the option with the highest Option Evaluation is selected: after the execution of P8 the atom

maximal_option_evaluation(O)

holds, where O is the option with the highest Option Evaluation Value.

P10 maximal_option_evaluation(O) $\twoheadrightarrow$
selected_option(O) & decision_process_ended

Property P10 specifies the last step in the Decision Making Process. It defines the output of the process and ensures that the atom

decision_process_ended

holds so that the persistence rules no longer apply to future states.

Adaptation of the Somatic Markers

As Somatic Marking is a process rooted in experience, the model includes a mechanism for adapting the Somatic Markers according to the evaluations of outcomes that result from the execution of the selected option. This mechanism consists of a update function that takes both previous and current experiences in account. This formula, explained in more detail in (Jonker and Treur, 1999), is as follows:

$$smv(O)_t = (1\text{-}d) \cdot smv(O)_{t\text{-}1} + d \cdot ev(O)_{t\text{-}1} \quad (1)$$

In this formula, the variable $smv(O)_t$ is the value of the Somatic Marker of option O at time t. The variable ev(O) is the evaluation value, a real value in the range [0,1], which depends on the evaluation function and the outcome of executing the selected option. The parameter d is a real value in the range [0,1] which determines the decay of the memory of previous experiences . A high value for d will cause the Somatic Markers to rapidly change in accordance with the evaluation values. In other words, the parameter d determines to what degree previous experiences are retained in relation to new experiences. A lower value for d will result in a more stable memory of experiences, while a higher value for d results in a Somatic Marker that is heavily influenced by recent experiences.

For example, at time t the Somatic Marker value of the option increase altitude has value 0.8. This option is executed and the outcome is evaluated as being good, i.e. ev(increase_altitude) = 1.0. The decay parameter d is set on 0.2. The updated value for the Somatic Marker for increase_altitude will be (1-0.2) · 0.8 + 0.2 ·1.0 = 0.84.

The following LEADSTO rules define the adaptation process:

P11 outcome_of_execution(OC) & belief_current_situation(S) &
selected_option(O) $\twoheadrightarrow$ evaluating_outcome_started(S, O, OC)

P12 evaluating _outcome_started(S, O, OC) &
outcome_evaluation(OC, EV) & somatic_marker(S, O, V) &
value(decay_parameter, D) $\twoheadrightarrow$
somatic_marker(S, O, (1-D)*V + D*EV)

P11 combines the outcome the agent receives from the environment with the current situation and selected option into a single atom start_evaluating_outcome(s, o, oc). P12 adapts the Somatic Marker associated with the relevant situation and option according to formula (1).

SIMULATION

Setup

The model described in the previous sections has been used to run a number of simulations, using the LEADSTO software environment as described in (Bosse et al, 2007). An environment and scenarion for the agent has been implemented based on the case described earlier.

The decay parameter and the initial Somatic Marker values have been varied between the simulations A total of 12 different configurations have been run, with 4 different values for the decay parameter (0.2, 0.4, 0.6 and 0.8) and three different assignments of initial Somatic Marker values.

The first assignment of initial Somatic Marker values consists of assigning the value 0.5 to each Somatic Marker. This value represents the neutral value, with higher values indicating a more positive feeling and lower values indicating a less positive feeling.

The second and third assignments of initial Somatic Marker values represent biases or personal preferences. They consist of assigning the neutral value 0.5 to each Somatic Marker except a subset of Markers which are set to the value 1.0. In the second assignment the options with the value 1.0 are chaff_and_beam and flares, in the third assignment the option reroute has the value 1.0.

The idea behind this setup is that the second assignment represents a ‘squandering’ agent and the second assignment a ‘cautious’ agent. At least at the start the ‘squandering’ agent has a preference for the options that require the most resources to execute. The ‘cautious’ agent has a preference at the start for the always save option reroute. In this way it can be tested if the model can represent certain personal styles of decision making.

Simulation Traces

The LEADSTO software environment has been used to generate a number of traces. For each of the twelve configurations four traces have been generated. In each trace one of the four possible situations (radar-guided SAM, IR SAM, MANPAD or AAA) is observed by the model of the agent 25 times. The number of states between each option execution is 50 steps, as it takes the model roughly that many states to proceed from observing a situation to executing an option. For all traces the threshold for Option Elimination has been set at 0.2.

Within this software environment simulation traces (i.e., sequences of states) can be visualised. An example of such a simulation trace can be seen in Figure 3. Here, time is on the horizontal axis, the state properties are on the vertical axis. A dark box on top of the line indicates that the property is true during that time period, and a lighter box below the line indicates that the property is false.

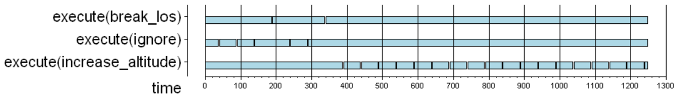

Figure 3: Partial trace showing option executions for the settings (situation AAA, d = 0.2, neutral initial somatic markers)

Figure 3 shows that the agent first selects the ignore option twice. The reason for this is that at the start all Somatic Evaluation Values are identical for all options and the rational utility of the ignore option is the highest. After two tries, the Somatic Marker value for the option ignore in the AAA situation has decreased, as can be seen in Figure 4.

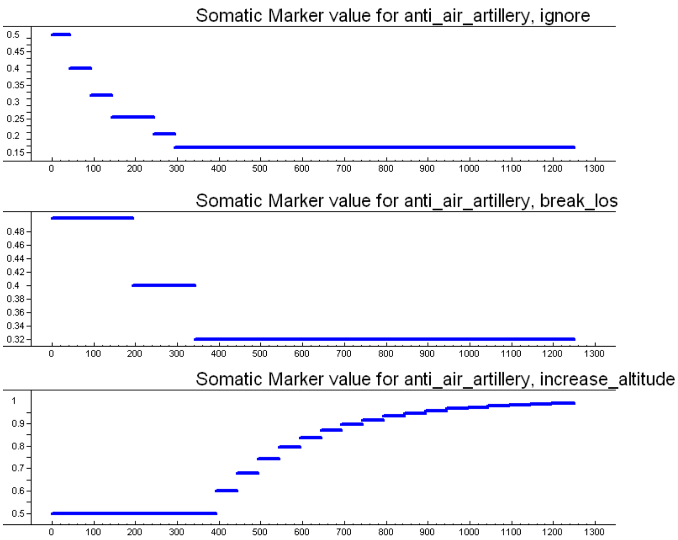

Figure 4: development of Somatic Marker values for the settings (situation AAA, d = 0.2, neutral initial Somatic Markers)

The option break_los also result in a low outcome evalution and this the corresponding Somatic Marker value becomes lower. Note that the option ignore will not be considered anymore, as its Somatic Evaluation Value (which in the current model is equal to the Somatic Marker value) falls below the threshold value of 0.2. After the eighth observation, the option break_los will no longer be selected but will still be considered in the Rational Analysis process as its Somatic Marker stays above the threshold at value 0.32.

As there are 48 traces, it is not feasible to show for all traces the option execution traces and Somatic Marker values development. In the next section a formal and automated way of verificating the traces is shown.

PROPERTIES VERIFIED UPON THE TRACES

In order to verify whether the simulation results comply to the desired behavior of the model, formal verification techniques have been used. First, the techniques are explained, followed by the actual verification of the properties.

The language used: Temporal Trace Language (TTL)

The verification of properties has been performed using a language called TTL (for Temporal Trace Language), cf. (Bosse et al., 2005) that features a dedicated editor and an automated checker. This predicate logical temporal language supports formal specification and analysis of dynamic properties, covering both qualitative and quantitative aspects. TTL is built on atoms referring to *states* of the world, *time points* and *traces*, i.e. trajectories of states over time. In addition, *dynamic properties* are temporal statements that can be formulated with respect to traces based on the state ontology Ont in the following manner. Given a trace γ over state ontology Ont, the state in γ at time point t is denoted by state(γ, t). These states can be related to state properties via the formally defined satisfaction relation denoted by the infix predicate |=, i.e., state(γ, t) |= p denotes that state property p holds in trace γ at time t. Based on these statements, dynamic properties can be formulated in a formal manner in a sorted first-order predicate logic, using quantifiers over time and traces and the usual first-order logical connectives such as ¬, ∧, ∨, ⇒, ∀, ∃. For more details on TTL, see (Bosse et al., 2005).

Verification Results

A number of properties have been identified to verify the behavior of the proposed model. A first property that has been identified specified that when an option is selected which turns out to be bad, the **Somatic Marker** value for that option is lowered:

DP1: Lowering Somatic Marker values
If an option O has been selected, and the evaluation of this option is bad, then the **Somatic Marker** value of this option will be lower than before.
Formally:
∀γ:TRACE, t1:TIME, O:OPTION, OC:OUTCOME, S:SITUATION, V1:REAL
[[state(γ, t1) |= selected_option(O) &
state(γ, t1) |= belief_current_situation(S) &
state(γ, t1) |= somatic_marker(S, O, V1) &
state(γ, t1) |= outcome_of_execution(OC) &
state(γ, t1) |= evaluation(OC, 0)]
⇒ ∃t2:TIME > t1, V2:REAL [state(γ, t2) |= somatic_marker (S, O, V2) & V2 < V1]]

This property is satisfied for all the traces that have been generated.

Given the fact that such **Somatic Markers** are lowered in value, the **Somatic Marker** values will eventually become lower than a the threshold value. Property DP2 expresses that options which have a **Somatic Marker** value below the threshold will be ignored:

DP2: Ignoring values below threshold

If the Somatic Marker value for an option O is below the threshold, then this option is never selected.

Formally:

∀γ:TRACE, t1:TIME, O:OPTION, V, TH:REAL
[[state(γ, t1) |=somatic_marker (S, O, P) &
state(γ, t1) |= value(threshold, TH) &
P < TH]
⇒ ¬∃t2:TIME > t1 [state(γ, t2) |= selected_option(O)]]

Again, this property is satisfied for all traces.

The combination of DP1 and DP2 leads to bad options no longer being selected within a certain duration, as specified in property DP3.

DP3(D:DURATION): Ruling out bad options

If an option O has been selected, and the evaluation of this option is bad, then within duration d this option will no longer be selected.

Formally:

∀γ:TRACE, t1:TIME, O:OPTION, OC:OUTCOME, S:SITUATION, V1:REAL
[[state(γ, t1) |= selected_option(O) &
state(γ, t1) |= current_situation(S) &
state(γ, t1) |= somatic_marker (S, O, V1) &
state(γ, t1) |= outcome_of_execution(OC) &
state(γ, t1) |= evaluation(OC, 0)]
⇒ ¬∃t2:TIME > t1 + D [state(γ, t2) |= selected_option(O)]]

The results of the verification with a duration setting of 100 time steps are shown in the table below.

Table 3. Result of verification for property DP3.

Situation	Decay parameter	Outcome (initial setting)
AAA	0.2	**not** satisfied (all initial settings)
AAA	0.4	satisfied (all initial settings)
AAA	0.6	satisfied (all initial settings)
AAA	0.8	satisfied (all initial settings)
IR SAM	0.2	**not** satisfied (all initial settings)
IR SAM	0.4	satisfied (all initial settings)
IR SAM	0.6	**not** satisfied (all initial settings)
IR SAM	0.8	satisfied (all initial settings)
MANPAD	0.2	satisfied (all initial settings)
MANPAD	0.4	satisfied (all initial settings)
MANPAD	0.6	satisfied (all initial settings)
MANPAD	0.8	satisfied (all initial settings)
Radar SAM	0.2	**not** satisfied (all initial settings)
Radar SAM	0.4	satisfied (all initial settings)
Radar SAM	0.6	**not** satisfied (all initial settings)
Radar SAM	0.8	satisfied (all initial settings)

As can be seen, the property is not satisfied in all traces. The expectation was that this property would be satisfied for the higher decay values whereas it is not satisfied for the lower values. As can be seen however, the property is satisfied for all traces with both decay 0.4 and 0.8, but is not satisfied for some of the traces with an decay setting of 0.2 and 0.6. Analysis of these traces has shown that in the particular circumstances, a bad option was selected once, followed by a selection of another bad option. Due to the fact that the latter option was also bad (and hence, the value for the option was lowered), the option with the highest value was again the former. This results in the option being selected again outside of the 100 time steps region from the first time selection. As a consequence, the property is not satisfied. For the case with decay 0.4 bad options are also selected twice but consecutively (within 100 time steps), resulting in the property being satisfied.

Besides the properties concerning the bad options being ruled out, the good options will be selected time after time once such a good option has been found.

DP4: Continue with good options

If an option O has been selected, and the evaluation of this option is good, then this option will be selected the next time as well.

Formally:

∀γ:TRACE, t1:TIME, O:OPTION, OC:OUTCOME, S:SITUATION, V1:REAL
[[state(γ, t1) |= selected_option(O) &
state(γ, t1) |= current_situation(S) &
state(γ, t1) |= somatic_marker (S, O, V1) &
state(γ, t1) |= outcome_of_execution(OC) &
state(γ, t1) |= evaluation(OC, 1)]
⇒ ∃t2:TIME > t1 [state(γ, t2) |= selected_option(O)]]

This property is satisfied for all traces.

DISCUSSION

Damasio's Somatic Marker Hypothesis (Damasio, 1994) shows how emotions play an essential role in decision making. It gives an account of how feeling (or experiencing) emotions in certain situations over time leads to the creation of a form of intuition (or experience) that can be exploited to obtain an efficient and effective decision making process for future situations met. Example of patients with brain damage related to feeling emotions show how inefficient and ineffective a decision making process can become without this somatic marking mechanism. Damasio's theory contrasts with the earlier tradition in decision making models, where the focus was on rational decision making based on the Expected Utility Theory, and where the aim was to exclude effects of emotions and biases on decision making; e.g., (von Neumann and Morgenstern, 1944; Friedman and Savage, 1948; Arrow, 1971; Keeney and Raiffa, 1976).

As for fighter pilots crucial decisions have to be made in very short times, it seems plausible that they heavily rely on such mechanisms. The computational model for pilot decision making described in this paper formalises the ideas informally described by Damasio (1994). When applied to specific scenarios, the model shows patterns as can be expected according to Damasio's theory.

Creating the model is one of the first steps in larger research programme. In next steps the model will be compared to decision making behaviour of human pilots in a simulation-based training setting. One of the questions that can be addressed is in how far the effect of a very intense or traumatic experience, for example a narrow escape, is covered well by the model. Another question that can be addressed is how the model can be used (or extended) to describe the effect of specific simulation-based training programmes, where a specific type of situations are offered and/or guidelines are put forward to handle such situations.

REFERENCES

Arrow, K.J., (1971). *Essays in the Theory of Risk-Bearing.* Chicago: Markham, 1971.

Bittner J.L., Busetta, P., Evertsz, R., Ritter, F.E., (2008). *CoJACK – Achieving Principled Behaviour Variation in a Moderated Cognitive Architecture.* Proceedings of the 17th Conference on Behavior Representation in Modeling and Simulation 08-BRIMS-025 Orlando, FL: U. of Central Florida.

Bechara, A. & Damasio, A. (2004) *The Somatic Marker Hypothesis a neural theory of economic decision.* Games and Economic Behavior, vol. 52, pp. 336-372.

Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J., Specification and Verification of Dynamics in Agent Models. *International Journal of Cooperative Information Systems.* In press, 2009.

Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J., (2008). A Language and Environment for Analysis of Dynamics by Simulation. *International Journal of Artificial Intelligence Tools*, vol. 16, 2007, pp. 435-464.

Damasio, A., (1994). *Descartes' Error: Emotion, Reason and the Human Brain*, Papermac, London.

Damasio, A., (1999). *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, Harcourt Brace, Orlando, Florida.

Friedman, M., L. J. Savage (1948). The Utility Analysisof Choices Involving Risks, *Journal of Political Economy,* 56 (1948), 279-304.

Jonker, C.M., and Treur, J., (1999). *Formal Analysis of Models for the Dynamics of Trust based on Experiences*. In: F.J. Garijo, M. Boman (eds.), Multi-Agent System Engineering, Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'99. Lecture Notes in AI, vol. 1647, Springer Verlag, Berlin, 1999, pp. 221-232.

Kahneman, D., A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, Vol. 47, 1979, pp. 263-292.

Keeney, R. L., Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* New York: Wiley, 1976.

Reiter, R (2001). *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems.* The MIT Press.

Tversky,A ., Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases, *Science*, 185 (1974), 1124-1 131.

von Neumann, J., Morgenstern, 0 . (1944). *Theory of Games and Economic Behavior.* Princeton: Princeton University Press, 1944.