

An Agent Model for Decision Making Based upon Experiences Applied in the Domain of Fighter Pilots

Mark Hoogendoorn¹, Robbert-Jan Merk^{1,2} and Jan Treur¹

¹VU University Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081a, Amsterdam, the Netherlands
{mhoogen, treur}@few.vu.nl

²National Aerospace Laboratory, Training, Simulation & Operator Performance
Anthony Fokkerweg 2 Amsterdam, the Netherlands
merkj@nlr.nl

Abstract — In this paper a computational model for decision making based on experiences is presented, inspired by the Somatic Marker Hypothesis. The use of the model is illustrated for the domain of fighter pilot decision making. Hereby, simulation runs have been performed upon this scenario, and the results thereof have been formally verified based upon properties inspired on Damasio’s Somatic Marker Hypothesis.

Keywords-Decision making, Somatic markers

I. INTRODUCTION

Decision making usually involves expectations about possible consequences of decision options and uncertainty about them. Traditionally the literature on decision making was dominated by the Expected Utility Theory [1] [10] [14] [17]. Here, decision making takes place by calculating 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 bad in estimating probabilities, and also may allow emotions and biases to play a role in a decision making process, as is found in several experiments; e.g., [13][16].

There are alternative theories on human decision making that do not assume an ideal rationalism and are more naturalistic. These theories claim that emotions, experience intuitive thinking and other non-rational aspects play a vital role in human decision making. Some of the more influential theories in the field of naturalistic decision making are the Recognition-Primed Decision model [14], the Cognitive Continuum Theory [11] and the Somatic Marker Hypothesis [7]. In this paper the focus is on the Somatic Marker Hypothesis.

In his research, Damasio 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 relate to sensed body states [7].

In this paper a computational decision making model is presented which draws inspiration from the Somatic Marker Hypothesis. The main purpose of this model is to create agents which show realistic humanlike behaviour. Because of this, the idealised rational approach to decision making is unsuitable. The Somatic Marker Hypothesis offers an interesting view on human decision making on which a computational model can be based in a relatively straightforward way.

While other computational models based on the Somatic Marker Hypothesis have been developed (see e.g. [5]), the innovation of this paper lies in the development of a computational model that is inspired by the Somatic Marker Hypothesis that can be used for decision making in agents.

Although the validity of the theory is sometimes doubted (see e.g. [9]), it can still be considered to be a useful source of inspiration for the development of agents for the aforementioned purpose. Moreover, the claims of validity of the computational model presented here are based on the results of simulations instead of solely relying on the validity of the Somatic Marker Hypothesis.

The use of the model is illustrated for the domain of fighter pilot decision making. First, the Somatic Marker Hypothesis is explained in more detail, after which the computational model is described. Thereafter, simulation results are presented, including formal properties that have been verified against the generated results.

II. DECISION MAKING AND EXPERIENCE

The Somatic Marker Hypothesis provides a theory on decision making which dedicates a central role to experienced emotions. 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.’ ([7] p. 173)

This theory consists of two main ideas:

- (1) the way in which somatic markers affect (future) decisions,
- (2) the way in which somatic markers depend on (past) experiences.

Concerning (1), if a decision is to be made between options which can lead to potentially harmful or advantageous outcomes, each of such options induces a somatic response which is experienced 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 that option. Damasio describes the use 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*’ ([7], 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’. ([7], p. 174)

Concerning (2), 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.’ ([7], 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.

III. MODEL DESCRIPTION

The model has been defined as a set of temporal relations between properties of states. A state property is a conjunction of atoms or negations of atoms that hold or do not hold at a certain time. 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, transitions between states are defined.

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 $\alpha \rightarrow_{e, f, g, h} \beta$, 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 .

A specification of dynamic properties in LEADSTO format has advantages that it is executable and that it can often easily be depicted graphically. LEADSTO has been used to model different kinds of cognitive processes with success (see [5] for some examples). For more details of the LEADSTO format, see [6]. As all of the temporal relations used in the model are of the form $\alpha \rightarrow_{0,0,1,1} \beta$, the notation $\alpha \rightarrow \beta$ will be used instead.

A. The Decision Making Process

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. The situation is represented by an atom supplied by the environment and can be seen as the result of the agent’s perception of its environment. For example, in the case described earlier, the agent could encounter an enemy fighter from its side. In the model the decision making process would receive the atom `observed(enemy_from_side)`.

In the Decision Making process for each option the option preference, a real number between 0 and 1, is determined. Both somatic markers and rational utility values are used to calculate option preferences. The option with the highest option preference is then selected for execution.

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 set of real numbers between 0 and 1, one per goal, where a higher value means a more positive evaluation. These evaluation values are then used to adapt the appropriate somatic markers associated with each goal. 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.

B. Step 1: Somatic Evaluation

The purpose of the Somatic Evaluation process is to assign a real value between 0 and 1 to each option. This value, the somatic evaluation value, is determined per option by adding the weighted values of the different types of somatic markers associated with the option and current situation. For each goal the agent has, there is a different type of somatic marker. There is also a weight value for each type of somatic marker with which the value of the somatic marker is multiplied. This way, it is possible to vary the influence each type of somatic marker has on the final somatic evaluation value, which can be used to represent personal characteristics. The formula for determining the somatic evaluation value is:

$$sev(O)_t = w_1 \cdot smv(G_1, O)_t + w_2 \cdot smv(G_2, O)_t + \dots + w_n \cdot smv(G_n, O)_t$$

where $sev(O)_t$ is the somatic evaluation value for option O at time t , $smv(G_i, O)_t$ the value for the somatic marker associated with goal G_i at time t , w_i the weight for goal G_i . Note that the somatic markers are those for the current situation. The weights add up to 1, so that the somatic evaluation value remains within 0 and 1. For the sake of brevity the temporal properties defining this process has been omitted.

C. Option Elimination

The next step is the Option Elimination process. All the atoms of the form $somatic_evaluation_value(O, V)$, generated in the Somatic Evaluation process, are transformed into atoms of the form $remaining_somatic_evaluation_value(O, V)$ if V does not fall below a threshold value. All other atoms of that form are discarded, effectively eliminating the options associated with those atoms. In P1 and P2 this process is formalised.

P1
 $somatic_evaluation_value(O, V)$ &
 $value(threshold, Th)$ &
 $V \geq Th$
 \rightarrow
 $remaining_somatic_evaluation_value(O, V)$ &
 $somatic_evaluation_ended$

P2
 $remaining_somatic_evaluation_value(O, V)$ &
 $not(decision_making_ended)$
 \rightarrow
 $remaining_somatic_evaluation_value(O, V)$

D. Rational Analysis

The next subprocess is the Rational Analysis. In this process a rational utility is calculated for each option for which the atom $remaining_somatic_evaluation(O, V)$ holds. According to [7], the rational phase is partly influenced by the preceding somatic marking. For this reason the assumption is that the remaining somatic markers are used in determining the outcome of the rational phase, which is a number of utility values.

In the design of the model there are atoms of the Form $belief(utility(S, O, U))$ which couple each situation S and each option O with a real value U between 0 and 1, indicating the utility for that option in that particular situation. More

elaborate utility functions are certainly possible but fall outside the scope of this paper.

P3
 $remaining_somatic_evaluation_value(O, V)$ &
 $belief(current_situation(S))$ &
 $belief(utility(S, O, U))$
 \rightarrow
 $option_utility(O, U)$

P4
 $remaining_somatic_evaluation_value(O, V)$ &
 $option_utility(O, U)$ &
 $value(rational_ratio, R)$
 \rightarrow
 $option_preference(O, R * U + (1-R) * V)$

Property P3 defines the determination of an option utility for each remaining option. This consists of attaching to each remaining option the utility that the agent believes is the expected utility for that option considering the current situation. In P4 for each remaining option the option preference is determined. This value is taken as a weighted average between the somatic evaluation value and the option utility. The parameter rational ratio determines what weight the option utility has in determining the option preference. In other words, a higher rational ratio shifts the Decision Making process more towards the rational side, while a lower rational ratio makes the Decision Making process more intuitive.

After P4 has been applied, the selected option is determined by taking the option with the highest option preference. The temporal properties that define this final selection are not included in this paper for the sake of brevity.

The selected option is then executed, which results in some outcome that is used for adapting the somatic markers.

E. 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. An update function described in [12] has been chosen to represent the Somatic Marker adaptation mechanism. This is a typical update function, similar to many other kinds of update function that are regularly used for modelling dynamics but it is certainly possible to use similar functions. The following formula describes the update function as used in the model:

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

In this formula, the variable $smv(G, O)_t$ is the value of the somatic marker of option O associated with goal G at time t . The variable $ev(G, O)_t$ is the evaluation value, a real value between 0 and 1. The parameter d is a real value also between 0 and 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.

Determining the evaluation value is based on the concept of a body state. In [7], p. 180, Damasio states that

‘At the neural level, somatic markers depend on learning within a system that can connect certain categories of entity with the enactment of a body state, pleasant or unpleasant.’ ([7], p. 180)

So it appears that the body either reacts positively or negatively in response to the outcome of an action. The precise dynamics of what body state is generated depends on innate dispositions (primarily survival related), and social conditioning. In the model this is represented by a number of atoms of the form $\text{resulting_body_state}(G, Oc, V)$, one for each goal-outcome combination.

The following LEADSTO rules show how the somatic marker adaptation is modelled.

P5

```
belief(outcome(Oc)) &
belief(current_situation(S)) &
belief(selected_option(O)) &
resulting_body_state(Oc, G, V)
→
evaluation(G, O, V)
```

P6

```
evaluation(G, O, Ev) &
somatic_marker(G, S, O, Smv) &
value(decay_parameter, D)
→
updated_somatic_marker(G, S, O) &
new_somatic_marker_value(G, S, O, (1-D) * Smv + D * Ev)
```

Each time the agent observes the outcome of an option it executed, it determines a resulting body state. In P5, the value v of the $\text{resulting_body_state}$ relevant to the current outcome of an executed option is used as evaluation value in the update function to modify each somatic marker belonging to that option. This way, the agent learns from its experiences.

IV. CASE STUDY

In order to test the model, a case has been constructed that represents a simplified environment from the domain of fighter airplane combat. In this case there is a single fighter, controlled by an agent, which is flying some kind of mission. Its goal is to arrive at his target. However, at some point it detects an enemy aircraft. This forces the agent to make a decision on how to deal with this situation, which is done by an implementation of the model described in this paper.

There are 3 different situations that the agent can encounter: the enemy approaches from the front, the side or from behind. In this case, the agent has four options to deal with these situations:

- The agent can continue its flight in order to reach his target.
- The agent can engage the enemy
- The agent can turn around and return to base.
- The agent can take an detour to its target, which requires it to fly over the enemy anti-air position.

The outcome of the execution of one of these options depends on the current situation and is probabilistic determined. For example. executing the option engage_enemy in the situation enemy_from_behind has a 30% chance of the agent being shot down, a 50% chance of the agent defeating the enemy and reaching the target and a 20% chance of defeating the enemy and being force to return to base. In the next two sections the choices for determining the utility values and resulting body states are explained.

The utility for each option in each situation that has been chosen for this case are shown in Table 1.

TABLE I. UTILITIES

		Situation		
		Enemy from side	Enemy from behind	Enemy from front
Option	Continue-mission-direct-route	1	1	0
	Continue-mission-detour	1	1	1
	Engage-enemy	0,5	0,5	0,5
	Return-to-base	0,5	0	1

The reasoning behind this allocation of utility values is that mission success and survival have a higher priority than defeating the enemy fighter. In general the agent has the orders to try to complete the mission and to avoid the enemy fighter and only to engage the enemy fighter if the opportunity to do so is good enough in its own ‘opinion’.

Therefore the continue-mission options have high utility values, except when the enemy comes in from the front. In that situation the option $\text{continue-mission-direct-route}$ has low utility, as survivability is important and the agent has to try to avoid the enemy fighter. The option engage-enemy has always a medium utility, as it is left to the agent’s discretion to choose whether to engage. The utility for return-to-base is heavily dependent on the enemy fighter’s angle of approach: if the enemy comes from the front, continuing the mission will be dangerous and so return-to-base is a good option, while if the enemy comes from behind, return-to-base is a bad option as the agent has the orders to avoid the enemy.

Table 2 shows the representation of the resulting body states for each outcome. A value of 1 represents a positive body state, a value of zero a negative body state. The body states are coupled to goals and the allocation of values is based on how good an outcome is for reaching that goal.

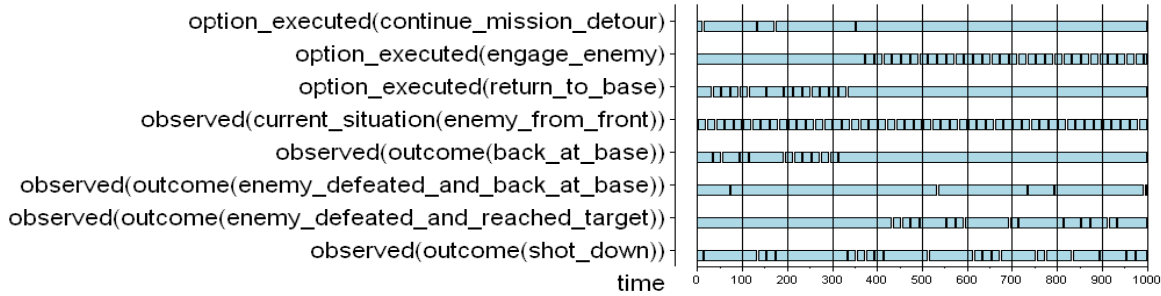


Figure 1. Events observed and options executed in situation enemy-from-front with weight setting 2.

TABLE II. RESULTING BODY STATES

		Goals		
		Lethality	Survivability	Resource control
Outcomes	Shot down	0	0	0
	Back at base	0	1	0
	Reached target	0	1	1
	Enemy defeated & reached target	1	1	1
	Enemy defeated & back at base	1	1	0

Lethality is about defeating the enemy, so all outcomes that include the defeat of the enemy result in a positive outcome. Being shot down is the only way of having a negative body state in regard to survivability, as in all other outcomes the agent survives unharmed. Finally, in this case resource control is mainly about fulfilling the mission objective, so all outcomes in which the target is reached result in a positive body state.

V. SIMULATION RESULTS

The model described in the previous sections has been used to run a number of simulations, using the LEADSTO software environment as described in [6]. An environment and scenario for the agent has been implemented based on the case described earlier. Hereby, all three scenarios as

presented before have been addressed.

In order to test whether different weights for somatic markers lead to different behaviour, for four different settings of somatic marker weights simulations have been run. The exact settings are shown in Table 3.

TABLE III. SOMATIC WEIGHT SETTINGS

Setting	W(Lethality)	W(Survivability)	W(resource control)
1	0,33	0,33	0,33
2	0,50	0,25	0,25
3	0,25	0,50	0,25
4	0,25	0,25	0,50

In setting 1, all types of somatic markers have equal influence in the determination of the somatic evaluation value. In settings 2, 3, and 4 the marker weights for respectively lethality, survivability and resource control are set higher, increasing the influence of the associated somatic markers on decision making.

Figures 1, 2 and 3 show the model's behaviour for the enemy-from-front scenario where setting 2 is used for the weight settings for the somatic markers within the model. Figure 1 gives an overview of the observations from the environment that the model is given during the simulation, the option the model executes and the outcome of executing that option.

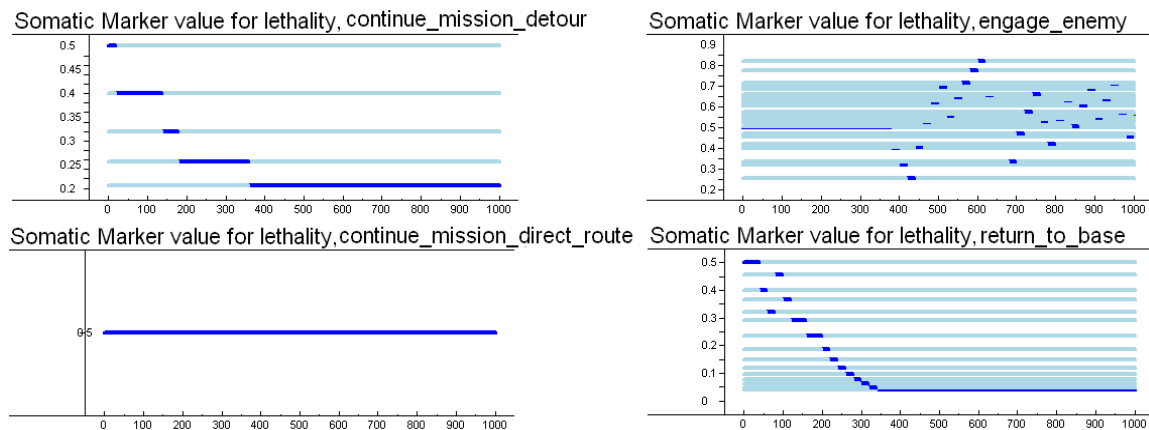


Figure 2. Value of the lethality somatic marker in situation enemy-from-front with weight setting 2.

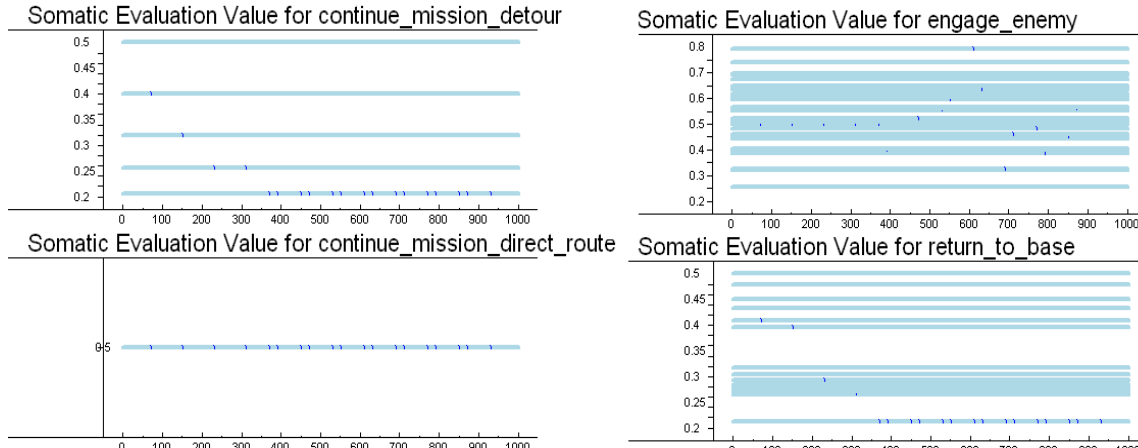


Figure 3. Change of somatic evaluation value over time with weight setting 2 in situation enemy-from-front.

Here, the atoms that occur during the simulation are shown on the y-axis whereas the x-axis shows a timeline. A dash on this timeline denotes the atom being true at that time point whereas a grey box indicates false. The figure shows that in the first iterations, the model selects return-to-base predominantly and that leads predominantly to the outcome back-at-base, with sometimes shot-down as outcome. Around iteration 14 (timestep 380), the model switches to executing engage-enemy, after which the outcome switches between enemy-defeated-and-reached-target and shot-down. As setting 2 represents an aggressive pilot who places a high value on lethality, this switch to a more offensive option is as expected.

Figure 2 shows the progression of the Somatic Marker value for lethality for the same situation. In this Figure, the x-axis again represents time, and the y-axis represents the value. The figure shows that the somatic marker's value drops for the options return-to-base and continue-mission-detour, as these options rarely to never lead to the defeat of the enemy. Later when the model starts to execute engage-enemy, the somatic marker starts to fluctuate, as the outcome for that option is not always the defeat of the enemy. However, the value is higher than neutral over time as the chance of defeating the enemy is higher than 50%.

Figure 3 shows the Somatic Evaluation Value for each option in the scenario. As the figure shows, at first the evaluation values are high for return-to-base and continue-mission-detour but drop as the outcomes are unfavourable. Later the selection of the engage-enemy option causes the evaluation value of that option to fluctuate and as the other options are no longer selected, the evaluation values of the other options to remain the same.

For every situation-weight setting combination, a simulation has been run. In each simulation the decision making model receives 50 times the same situation to decide on. The results of these simulations have been verified, as shown in the next section.

Table 4 shows how many times each option has been selected with different somatic weight settings for the enemy-from-front scenario.

TABLE IV. OPTION SELECTION IN SITUATION ENEMY-FROM-FRONT

		Somatic weight setting			
		Setting 1	Setting 2	Setting 3	Setting 4
Option	Continue-mission-direct-route	49	48	49	49
	Continue-mission-detour	1	2	1	1
	Engage-enemy	0	0	0	0
	Return-to-base	0	0	0	0

In this situation, when the somatic markers associated with the lethality goals have a higher weight, the option engage-enemy is selected much more often than with a neutral setting. This also the case to a lesser extent when resource control has a higher weight, as in this situation the option engage-enemy leads much more often to the outcome reached-target than any other option. There is little difference between the results of the neutral setting and setting 3, where survivability has a higher weight, as in the neutral setting return-to-base is already predominantly chosen. This is probably due to the allocation of utility values, in which a high emphasis is laid upon survivability.

In Table 5 and 6 the option selection for the other two situations are shown.

TABLE V. OPTION SELECTION IN SITUATION ENEMY-FROM-SIDE

		Somatic weight setting			
		Setting 1	Setting 2	Setting 3	Setting 4
Option	Continue-mission-direct-route	11	36	47	48
	Continue-mission-detour	4	4	3	2
	Engage-enemy	19	9	0	0
	Return-to-base	13	1	0	0

TABLE VI. OPTION SELECTION IN SITUATION ENEMY-FROM-BEHIND

		Somatic weight setting			
		Setting 1	Setting 2	Setting 3	Setting 4
Option	Continue-mission-direct-route	0	0	0	0
	Continue-mission-detour	3	4	2	4
	Engage-enemy	0	32	0	13
	Return-to-base	47	14	48	33

The somatic evaluation value for the option continue-mission-direct-route does not change, as this option is never selected. The somatic evaluation value for continue-mission-detour drops under the threshold of 0.25 after 4 selections, which means that this option will not be considered again and consequently not be selected at all. For the option return-to-base this happens after 14 selections. The somatic evaluation value for engage-enemy fluctuates strongly as there is a great variation between differing outcomes which lead to different resulting body state values.

This example shows that the agent has learned that in this situation continue-mission-detour and return-to-base are bad options and will only consider engage-enemy and continue-mission-direct-route in the future.

VI. VERIFICATION

In order to verify whether the behaviour of the model indeed complies to the Somatic Marker Hypothesis as proposed by Damasio, a logical verification tool has been used. Below, the formal language underlying this verification tool is explained, after which properties are shown that have been verified against a variety of traces.

The verification of properties has been performed using a language called TTL (for Temporal Trace Language), cf. [4] 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 $\text{state}(\gamma, t)$. These states can be related to state properties via the infix predicate \models , where $\text{state}(\gamma, t) \models p$ denotes that state property p holds in trace γ at time t . Based on these statements, dynamic properties can be formulated in a sorted first-order predicate logic, using quantifiers over time and traces and the usual first-order logical connectives such as \neg , \wedge , \vee , \Rightarrow , \forall , \exists . For more details, see [4].

The properties that have been verified against the simulation traces are shown below.

The first property (P1) expresses that a negative evaluation of an option in a given situation with respect to a certain goal results in the somatic marker value for that option going down.

P1: Lowering specific somatic marker value

If an option O has been selected, and the evaluation of this option with respect to a goal G is bad, then the somatic marker value of this option for goal G will be lower than before.

$$\begin{aligned} &\forall \gamma:\text{TRACE}, t1:\text{TIME}, O:\text{OPTION}, S:\text{SITUATION}, \\ &G:\text{GOAL}, V1:\text{REAL}, E:\text{REAL} \\ &[[\text{state}(\gamma, t1) \models \text{belief}(\text{selected_option}(O)) \ \& \\ &\quad \text{state}(\gamma, t1) \models \text{belief}(\text{current_situation}(S)) \ \& \\ &\quad \text{state}(\gamma, t1) \models \text{somatic_marker}(G, S, O, V1) \ \& \\ &\quad \text{state}(\gamma, t1) \models \text{evaluation}(G, O, E) \ \& \ E < \text{NEUTRAL}] \\ &\Rightarrow \exists t2:\text{TIME} > t1, V2:\text{REAL} \\ &\quad [\text{state}(\gamma, t2) \models \text{somatic_marker}(G, S, O, V2) \ \& \ V2 < V1]] \end{aligned}$$

In case the overall evaluation of an option in a given situation is below neutral, then the total somatic evaluation value goes down. This is expressed in property P2. The overall evaluation value is the weighted sum of the evaluation values for all goals. Note that the remaining formal forms have been omitted for the sake of brevity.

P2: Lowering overall evaluation value

If an option O has been selected, and the overall evaluation of this option is bad, then the value of the total somatic evaluation value for option O will go down.

The idea of Damasio is that certain options are no longer considered because they are not appropriate in a given situation. This idea is expressed in property P3 which states that once the total somatic evaluation value is below the threshold, the option will no longer be selected.

P3: Ignoring values below threshold

If the total somatic evaluation value for an option O is below the threshold, then this option is never selected.

Finally, property P4 expresses that eventually an option is selected which has a higher evaluation value than neutral.

P4: Eventually a good option is selected

There exists a time point such that an option O is selected which scores good for all goals.

The properties above have been verified against 12 simulation traces (3 situations, each consisting of 4 settings) During the verification process, a value of 0.5 has been used for the constant NEUTRAL. It was shown that property P1-P3 are satisfied for all traces. Property P4 however is not satisfied for the case whereby the enemy comes from the front, and the weight setting 3. The same holds for the case enemy from behind with setting 2. This is due to the fact that the probability of an option having a positive evaluation for these scenarios is very small, and does not occur in the trace which has been checked.

VII. CONCLUSIONS

Damasio's Somatic Marker Hypothesis [7] 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. Examples 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., [1][10][14][17].

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

- For a given type of emotion, somatic markers are related to combinations of contexts and decision options for this context.
- When a decision has to be made within a given context, somatic evaluation values associated to the options 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.
- Within a given context, every decision option gets a somatic evaluation value associated based on the somatic markers.
- Decision options with low associated somatic evaluation value are eliminated from further decision processing.
- For the remaining decision options a (utility-based) rational analysis is made in which the somatic evaluation values serve as biases.
- Based on experiences for outcomes of chosen options for a given context, the somatic markers are adapted over time.

As for fighter pilots crucial decisions have to be made in very short times, it seems plausible that they heavily rely on such mechanisms. 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 program. In next steps the model will be compared to decision making behaviour of human pilots in a simulation-based training setting.

ACKNOWLEDGMENTS

This research has been conducted as part of the "SmartBandits" PhD project which is funded by the National Aerospace Laboratory (NLR) in the Netherlands. Furthermore, the authors would like to thank Jan Joris Roessingh, and the anonymous reviewers for their useful comments that helped to improve the paper.

REFERENCES

- [1] Arrow, K.J., (1971). *Essays in the Theory of Risk-Bearing*. Chicago: Markham, 1971.
- [2] 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.
- [3] Bechara, A. & Damasio, A. (2004) *The Somatic Marker Hypothesis a neural theory of economic decision*. Games and Economic Behavior, vol. 52, pp. 336-372.
- [4] Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J., (2009). Specification and Verification of Dynamics in Agent Models. *International Journal of Cooperative Information Systems*, vol. 18, 2009, pp. 167 - 193..
- [5] Bosse, T., Jonker, C.M., and Treur, J., (2008). Formalisation of Damasio's Theory of Emotion, Feeling and Core Consciousness. *Consciousness and Cognition Journal*, vol. 17, 2008, pp. 94–113.
- [6] 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.
- [7] Damasio, A., (1994). *Descartes' Error: Emotion, Reason and the Human Brain*. Papermac, London.
- [8] Damasio, A., (1999). *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, Harcourt Brace, Orlando, Florida.
- [9] Dunn, B.D., Dalgleish, T., Lawrence A.D. (2005). The somatic marker hypothesis: A critical evaluation, *Neuroscience and Biobehavioral Reviews* 30 (2005) 239–271.
- [10] Friedman, M., L. J. Savage (1948). The Utility Analysis of Choices Involving Risks, *Journal of Political Economy*, 56 (1948), 279-304.
- [11] Hammond, K. R., Hamm, R. M., Grassia, J., & Pearson, T. (1987). *Direct comparison of the efficacy of intuitive and analytic cognition in expert judgment*. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC- 17(5), 753-770.
- [12] 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*, Proc. 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.
- [13] Kahneman, D., A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, Vol. 47, 1979, pp. 263-292.
- [14] Keeney, R. L., Raiffa, H. (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: Wiley, 1976.
- [15] Klein, G.A., Calderwood, R. (1991). Decision models: some lessons from the field. *Systems, Man and Cybernetics*, *IEEE Transactions on Sep/Oct 1991* Vol. 21, 5, 1018-1026.
- [16] Tversky, A., Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases, *Science*, 185 (1974), 1124-1131.
- [17] von Neumann, J., Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944