

Towards a Tractable Appraisal-Based Architecture for Situated Cognizers

Alexander Staller and Paolo Petta
Austrian Research Institute for Artificial Intelligence*
Schottengasse 3, A-1010 Vienna, Austria (EU)
{alexs,paolo}@ai.univie.ac.at

**Extended version of the paper in: Canamero D., Numaoka C., Petta P. (eds.):
*Grounding Emotions in Adaptive Systems, Workshop Notes, Fifth International
Conference of the Society for Adaptive Behaviour (SAB98), Zurich, Switzer-
land, August 21, 1998.***

Abstract

This paper introduces TABASCO¹, an architecture for software agents aimed at integrating results from functional theories in emotion research and insights on the impact of the capacities and limitations of perception in a framework orientated along the situated “New AI”/ALife approach. This expository paper first briefly summarizes current views on the nature and function of emotion and then discusses related current appraisal theories in more detail. A survey of existing approaches to emotion synthesis is followed by a first outline of the TABASCO architecture, relating it to the areas of research in psychology, ALife and agent architectures.

1 Introduction

After having long been dismissed as irrational and of no utility, the Emotional is being rehabilitated as the system that continuously monitors the relationship of the individual to its world and instigates tendencies to act. Emotions are seen as a key element for successful coping with nondeterministic, highly dynamic environments, providing a flexible and adaptive mechanism. At the same time, “New AI”

*Supported by the Austrian Federal Ministry of Science and Transport.

¹The name TABASCO is derived from the project’s name “A Tractable Appraisal-Based Architecture for Situated Cognizers”.

aims at modelling adaptive behaviour in autonomous agents situated in complex, dynamic environments. It is therefore a reasonable goal to integrate a computational model of the emotion process within the architecture of a situated agent. This paper reports the first steps towards this end.

The paper is structured as follows. The next section sets out with a brief summary of the nature and adaptive function of emotion according to widely accepted theories and then moves on to a more detailed discussion of appraisal theories of emotion, which are gathering a strong and still increasing consensus in psychological research. Section 3 provides a review of research on agent architectures and emotion synthesis and presents a typical example of an existing appraisal-based architecture. Section 4 gives an outline of TABASCO, relating the architecture to the areas of research in psychology, ALife and agent architectures. Section 5 concludes the paper.

2 The Nature and Adaptive Function of Emotion

2.1 Emotion and Adaptation

Emotion can be viewed as a flexible adaptation mechanism that has evolved from more rigid adaptational systems such as reflexes and physiological drives ([38], [57], [64]). The flexibility of emotion is obtained by decoupling the behavioural reaction from the stimulus event. The heart of the emotion process thus is not a reflex-like stimulus-response pattern, but rather the *appraisal* of an event with respect to its adaptational significance for the individual, followed by the generation of an *action tendency* aimed at changing the relationship between the individual and the environment [28]. Lazarus [38] subsumes adaptive activities under the term *coping*. In the remainder of this section, appraisal, action tendencies, and coping are discussed in more detail.

2.2 Appraisal Theories of Emotion

“Appraisal theories” of emotion emphasize the role of a hypothesized continuous evaluation of the environment according to dimensions regarding the individual. This evaluation is called *appraisal*. It is conceived of as constituting element in emotion generation, mediating between events and emotions, explaining why the same event can give rise to different emotions in different individuals, or even in one and the same individual at different times. Conversely, appraisals offer a framework for the identification of the conditions for the elicitation of different emotions, as well as for understanding what differentiates emotions from each other.

Many theories have been developed in order to specify how many and which appraisal criteria are minimally needed for emotion differentiation (e.g. [28], [38], [47], [55], [57], [62]). There is a high degree of consensus with respect to these appraisal criteria (see e.g. [58] for a review). According to van Reekum and Scherer ([53], pp.259-260), these include “the perception of a change in the environment that captures the subject’s attention (novelty and expectancy), the perceived pleasantness or unpleasantness of the stimulus or event (valence), the importance of the stimulus or event to one’s goals or concerns (relevance and goal conduciveness or motive consistency), the notion of who or what caused the event (agency or responsibility), the estimated ability to deal with the event and its consequences (perceived control, power or coping potential), and the evaluation of one’s own actions in relation to moral standards or social norms (legitimacy), and one’s self-ideal”.

2.3 Modelling the Appraisal Process

The main application of the appraisal theories cited above is the structural analysis of the semantics of emotion words. However, these structural theories do not specify the nature of the appraisal process. In response to Zajonc [70], who criticized the “exaggerated cognitivism” of appraisal theories, several appraisal theorists pointed out that the appraisal process is not necessarily conscious or voluntary ([29], [57], [64]).

Leventhal and Scherer [41] attempt a more precise specification of the nature of the appraisal process. They suggest a hierarchical processing system consisting of three levels: sensory-motor, schematic, and conceptual. The sensory-motor level is based on innate hard-wired feature detectors giving rise to reflex-like reactions. The schematic level is based on schema matching. The schema concept was introduced by Bartlett [6] and has been widely used in cognitive science as a central memory structure. The conceptual level involves reasoning and inference processes that are abstract, active, and reflective. Problem-solving is an example of a conceptual-level process.

Smith et al. [63] suggest a model of the appraisal process that builds on the distinction between schematic and conceptual processing suggested by Leventhal and Scherer [41]. Schematic processing is fast, automatic, parallel, inflexible, and concrete. It can be thought of in terms of priming and spreading activation. In contrast, conceptual processing is slow, voluntary, serial, flexible, and relying on semantically accessible information. Smith et al. [63] especially emphasize the interactions between schematic and conceptual processing. For example, conceptual processing can activate, create, and alter schematic memories. On the other hand, a sufficiently activated schema can become available for conceptual processing. A novel feature of the model proposed by Smith et al. [63] is the so-called appraisal

register. It monitors appraisal information generated through schematic or conceptual processing, in addition to perceptual information, and generates an emotional reaction. The appraisal register is assumed to model the function of the amygdala, which plays an important role in the elicitation of fear [39].

Although the proposals by Leventhal and Scherer [41] and Smith et al. [63] are preliminary and not very detailed, results of research in the areas of social cognition, clinical psychology, and neuropsychology strengthen the evidence for a multi-level appraisal process [53].

2.4 Action Tendencies and Coping

The appraisal of the situation is just the first step of the emotion process. For successful adaptation, emotion must have an effect on the actions of the individual. According to Frijda [28], a “change in action readiness” is the essence of an emotion. As mentioned in section 2.1, emotion does not lead to a fixed action, but to the generation of an *action tendency*. Action tendencies “are states of readiness to achieve or maintain a given kind of relationship with the environment. They can be conceived of as plans or programs to achieve such ends, which are put in a state of readiness.” ([28], p. 75). Avoidance is an example for an action tendency during fear. With respect to action control, action tendencies have the feature of “control precedence”. However, whether an action tendency actually leads to action depends on feasibility tests and the monitoring of progress.

Lazarus [38] stresses that especially in humans action tendencies do not automatically lead to action. He emphasizes the importance of *coping*, which “consists of cognitive and behavioral efforts to manage specific external or internal demands (and conflicts between them) that are appraised as taxing or exceeding the resources of the person” ([38], p. 112). According to Lazarus, an action tendency is an innate biological impulse, while “coping is a much more complex, deliberate, and often planful psychological process” that “draws heavily on appraisals about what is possible, likely to be effective in the specific context, and compatible with social and personal standards of conduct (e.g., display rules, action rules)” ([38], p. 114). Coping processes can augment or inhibit the innate action tendencies.

To summarize, although Lazarus [38] agrees with Frijda [28] that flexible action tendencies are characteristic for emotion, he does not conceive of them as plans, in contrast to Frijda’s definition of action tendencies cited above. For Lazarus, plans are involved in coping, not in the generation of action tendencies.

3 Research on Agent Architectures and Emotion Synthesis

3.1 Agent Architectures

Wooldridge and Jennings [67] as well as Jennings, Sycara, and Wooldridge [37] provide a thorough review of existing agent architectures. Here only the distinction between *deliberative*, *reactive*, and *hybrid* architectures is briefly highlighted.

Wooldridge and Jennings ([67], p. 24) define a *deliberative* agent architecture “to be one that contains an explicitly represented, symbolic model of the world, and in which decisions (for example about what actions to perform) are made via logical (or at least pseudo-logical) reasoning, based on pattern matching and symbolic manipulation”. The characterizing feature of a deliberative agent is the ability to plan. Pfeifer [50] reviews the main problems of the symbol-processing approach (e.g. the well-known symbol grounding and frame problems).

These problems, together with theoretical results by Chapman [14] about the unfeasibility of the planning approach in practical applications, led to the emergence of “New AI” (“behaviour-based AI”, “situated AI”). Brooks [10, 11, 12], one of the main proponents of this new paradigm, argued that intelligence is not a property of disembodied systems, but a result of the interaction of an embodied agent with its environment. Further, he designed the subsumption architecture that does not make use of symbolic representations or reasoning at all. Rather, the architecture is a collection of task accomplishing behaviours that are implemented as augmented finite state machines arranged in a hierarchy. Higher layers represent more abstract behaviours and can be suppressed by the output of lower layers. The subsumption architecture is the prototypical example of a *reactive* agent architecture.

Despite the merits of reactive architectures, Jennings, Sycara, and Wooldridge [37] enumerate a number of disadvantages. For example, the fact that overall behaviour emerges from the interaction of simple behaviours, implies that there is no principled methodology for engineering agents to fulfill specific tasks. As a result, so-called *hybrid* architectures have been developed that combine aspects of both deliberative and reactive architectures. Typically, a hybrid architecture consists of a number of layers that are arranged hierarchically. Very common is a hierarchy of three layers. For example, the $3T$ architecture [9] consists of the layers “reactive skills”, “sequencing”, and “deliberation”. The reactive skills are coordinated by a skill manager [69]. At the sequencing layer, the Reactive Action Packages (RAPs) system [26] is used for activating and deactivating “sets of skills to create networks that change the state of the world and accomplish specific tasks” ([9], p. 238). The deliberation layer consists of the Adversarial Planner (AP) [23]. Other examples of hybrid architectures are TOURINGMACHINES [24, 25] and Cypress [66], which

combines a planner with a modification of the Procedural Reasoning System (PRS) [31].

3.2 Artificial Life: Situated Emotions

In a survey of earlier artificial intelligence research on models of emotion by Pfeifer [48, 49] the majority of efforts were found to be fundamentally hampered by the attempt to isolate the object of study according to what was been termed “the traditional research methodology”. Similar criticism had also already been voiced as early as in the late 70s by Aaron Sloman [59, 61]. Pfeifer’s review has had a seminal effect insofar as it substantially contributed to setting off a whole avalanche of related work in the Artificial Life community. The basic tenet of the Animat/Artificial Life approach is to thoroughly explore the space of opportunities offered by the most simple designs by virtue of their situatedness and physical embodiment, before moving up towards the incorporation of more “complex” components, such as declarative representations and related manipulation capabilities.

Nowadays, it is commonly recognized that research on functional process models of emotion synthesis requires *whole agent designs* ([51], [60]). Various approaches have been adopted to this end, adopting guidelines from ethology, neurophysiology, cognitive and social science ([4], [7, 8], [13], [33, 56], [42, 43, 54], [65], [68]). In the following, we briefly discuss the Affective Reasoner, a representative appraisal-based architecture.

3.3 An Existing Appraisal-Based Architecture

The Affective Reasoner [19, 20, 22] is an architecture for abstract, domain-independent emotion reasoning. It uses a descriptive process model of emotion based on a cognitive theory of emotion eliciting conditions [46, 47] that defines 24 discrete emotion types. The “rudimentary personalities” engendered by the Affective Reasoner architecture comprise an “interpretive personality component”—providing individuality with respect to their interpretation of situations—and a “manifestative personality component”—conferring individuality with respect to the way agents express or manifest their emotions. An agent’s interpretive personality component is defined by means of a hierarchical “goals, standards, and preferences” (GSP) database; its manifestative personality component is encoded in an action database.

Processing flows from some initiating event in a simulated world through emotion and action generation to the final stages simulating observation by other agents. Initiating situations are “construed” to fall into one of the eliciting condition classes: goal-relevant events; acts of accountable agents; attractive or unattractive objects; or combinations of the first three categories. In this way, each successful construal

instantiates an Emotion Eliciting Condition relation structure (EEC), a nine-tuple encoding of the characteristics of the situation relevant for further processing derived from the 6-slot structure employed in the BORIS system [17]. As the Affective Reasoner covers only a finite set of discrete emotions, the mapping of EECs to emotions involves some ad-hoc solutions to cope with a number of problems caused by the fact that often this mapping would not be unique but may reference multiple, and at times conflicting emotions.

Instantiated emotions are mapped to actions by means of an emotion manifestation lexicon. For each of the 24 emotion classes, 17 action response categories are defined, comprising: somatic responses; behavioural responses; communicative responses; evaluative responses; and cognitive responses such as initiation of plans. For deployment in a given domain, a few (about three, in the example given in the PhD thesis [19]) actual response actions (ranked by “intensity”) are to be defined for each action response category. The lexicon thus comprises over a thousand individual actions to be defined for each application domain. In addition, conflict sets have to be defined to explicitly rule out incompatible manifestations which may be triggered by co-occurring emotions.

The Affective Reasoner also includes support for social emotions, by virtue of a “Concerns of Other” database which at first is initialized with a partial, system-wide GSP database and extended via case-based abductive reasoning during the runtime of the program. Modeling of other agents is limited to at most “twice-removed” cases, i.e., the representation of another agents’ representation of the concerns of those agents which are important to it.

3.4 Discussion

The following is a summary of some noteworthy points:

- There exist robust, widely adopted solutions to fundamental aspects of the realization of situated software agents in virtual environments, in terms of hybrid layered architectures and organization of behavioural components, substantially facilitating the overall task of implementing such systems.
- Demonstrating the validity of Pfeifer’s view (section 3.2), abstract, domain-independent architectures such as the Affective Reasoner run into serious problems when deployed in interactive virtual scenarios, as also openly recognized by Clark Elliott [21]: to be effective in such applications, affective reasoning has to have appropriate access to pertinent information about and from the world, and has to be able to influence the overt external behaviour as well as the internal information processing of an agent. The only means

to achieve this is to fully integrate emotional competence into an architecture which in turn has to be adapted to the environment in which the agent is situated.

- Shallow approaches with a reified representation of a finite number of discrete emotional states, such as the Affective Reasoner, demonstrate the problems of brittleness and consistency well known from the traditional research area of expert systems.
- Reification of emotions as identifiable system components and routing of all processing through these entities engenders the problem of how to proceed from these emotions for further system processing, leading to the adoption of ad-hoc constructions of dubious validity.
- Discretizing emotions all but rules out smooth and fine-grained system behaviour; it is precisely this kind of capability that is becoming of increasing importance, as the information infrastructure in general, and virtual environments in particular increasingly assume the characteristics of persistent social places, in which long-term relationships develop.
- Applications of functional emotion theory can provide an alternative approach for tackling at least some of the problems mentioned, in good part by virtue of seeing emotions as processes, as characteristics of a whole architecture, which pervasively influence the overall behaviour of the agent.

TABASCO is an architecture for software agents situated in virtual environments that takes these points into account. It will be described in the following section.

4 The TABASCO Architecture

4.1 Outline

The ultimate goal of our research is the integration of the emotion process within the architecture of a situated software agent. We model emotions not as reified entities, but as adaptive processes related to the interaction of an agent with its environment. Especially, we need to model the steps that were identified in section 2 as characteristic of the emotion process: the appraisal process and the generation of action tendencies as well as coping activities.

The aim of this section is to provide a rough first outline of TABASCO. Despite its lack of detail, it should set the stage for further discussions and refinements. The agent architecture is shown in Figure 1 and will be explained in the following.

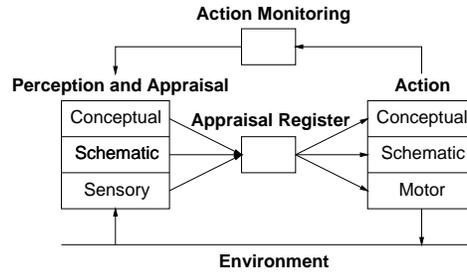


Figure 1: The TABASCO Architecture

The basic idea behind TABASCO is that the distinction between sensory-motor, schematic, and conceptual processing does not only apply to appraisal, but also to the generation of action, as proposed originally by Leventhal in his “perceptual-motor theory of emotion” (see e.g. [40]). So the two main components of the agent architecture **Perception and Appraisal** as well as **Action** are three-level hierarchies, with sensory processing as the lowest level of the **Perception and Appraisal** component and motor mechanisms as the lowest level of the **Action** component.

In the following, the components of TABASCO are described in more detail:

- **Perception and Appraisal**

This component processes environmental stimuli and models the appraisal process. The input must be evaluated according to the appraisal criteria enumerated in section 2.2. Frijda [29] argues that it is not necessary to model all appraisal checks explicitly. Rather, most of the appraisals “can be seen as implied by the very perceptual processes by which stimulus events are taken in” ([29], p. 380). For example, it is not necessary to model a “novelty check” explicitly. Rather, the appraisal of novelty is implied by the fact that the input does not match any schema.

Leventhal and Scherer [41] propose that at each level of processing the input is evaluated with respect to all appraisal criteria. But they do not provide empirical evidence for this claim. In our view, it may well be that an appraisal occurs not at all levels. For example, it is difficult to see how sensory processing is related to standards or norms.

The implementation of the several evaluations made during appraisal is still an open question. But research by social psychologists, especially in the field of social cognition, may help to model the processes involved in these evaluations (see [27] for an overview of research on social cognition). For

example, causal attribution has been investigated by social psychologists for decades. Findings of attribution theorists are relevant for the appraisal of agency or responsibility. Another research topic is the influence of social schemata on encoding, memory, and inference in social situations. The self is also an important research topic in social psychology. Self schemata have been proposed as representations of important attributes of the self. Further, research on self-regulation is relevant for modelling the appraisal process, e.g. for modelling the evaluation of one's coping potential or the comparison of oneself to standards and norms.

The processes at each level can be summarized as follows:

The sensory level consists of feature detectors for the detection of e.g. sudden, intense stimulation or the pleasantness of a stimulus.

At the schematic level, the input is matched with schemata, especially with social and self schemata. Smith et al. [63] propose the implementation of the schematic level as an associative network with a spreading activation mechanism. There are a number of associative network models of memory (e.g. ACT* by Anderson [3]). In addition, work on semantic networks in the fields of AI and information retrieval can guide the design of the schematic level.

The conceptual level involves abstract reasoning and inference based on propositional knowledge and beliefs. The inferences involved in causal attribution will be modelled at this level as well as the evaluation of one's actions in relation to norms or one's self-ideal. Therefore, models of self and others will be placed at the conceptual level. Ferguson's [24, 25] use of Belief-Desire-Intention (BDI) models at the "Modelling Layer" of his hybrid TOURINGMACHINES architecture as well as Elliott's "Concerns of Other" database (see section 3.3) are relevant for the design of the conceptual level.

- **Action**

We plan to design and implement this component by using an already existing hybrid agent architecture. A suitable candidate may be the 3T architecture [9] mentioned in section 3.1. The conceptual level will basically consist of a planner. According to Lazarus' [38] characterization of coping as a planful process, coping will be implemented at the conceptual level.

The 3T sequencing layer based on the RAP system [26] corresponds to schematic processing. In fact, Earl and Firby [18] merge ideas from the RAP system with Drescher's schema mechanism [16]. Action tendencies may be implemented using the RAP system. Although Frijda [28] uses the

term “plan” to characterize an action tendency, he also uses the term “flexible program” and states: “Flexible programs are those that are composed of alternative courses of action, that allow for variations in circumstances and for feedback from actions executed.” ([28], p. 83). This notion of a flexible program is very close to the functionality of a RAP, which is designed to react to feedback and “is simply a description of how to accomplish a task in the world under a variety of circumstances using discrete steps” ([9], p. 240).

The reactive skills of the $\mathcal{3}T$ architecture correspond to sensory-motor processing. In Figure 1 the lowest level of the **Action** component actually contains only the motor commands of the reactive skills, while the sensory input is processed at the lowest level of the **Perception and Appraisal** component. One important class of motor mechanisms relevant for modelling the emotion process are “expressive” motor mechanisms, such as facial and body expressions. These expressive behaviours can be implemented at the motor level of the **Action** component.

- **Appraisal Register**

This component mediates between the **Perception and Appraisal** and the **Action** components. The notion of an appraisal register is taken from Smith et al. [63]. It detects and combines the appraisal outcomes from the sensory, schematic, and conceptual levels of the **Perception and Action** component. Then it influences the **Action** component based on the appraised state of the world. This influence may have several forms: First, it may lead to the immediate execution of a motor program, such as a startle or an orienting response. These responses may not be considered as full emotions, but as “pre-emotions” [38] that enable the agent to gather more information for the appraisal process. Second, an action tendency may be generated by putting a RAP into the state of readiness. Third, a long-term planning process may be initiated in order to cope with the appraised event.

It should be noted that it is not necessary to wait for the appraisal outcomes of all levels of processing in order to influence the **Action** component. A startle response occurs immediately after the sensory input has been processed, not after slow conceptual processes. It is therefore a matter of discussion, whether the appraisal register is actually necessary. It may be possible to design a three-level architecture in which at each level of processing the appraisal process and the influence on action are directly coupled. On the other hand, the appraisal register has the advantage that the interplay between coping, action tendencies, and expressive behaviour can be better coordinated.

- **Action Monitoring**

Frijda ([29], p.381) notes that a “major source of appraisals is the monitoring of action planning and action execution”. For example, the exhaustion of the whole repertoire of plans and actions without success leads to the appraisal that no coping potential is available. Therefore, TABASCO contains an **Action Monitoring** component that monitors the planning and execution processes in the **Action** component and sends the results of the monitoring process to the **Perception and Appraisal** component, where it is integrated within the appraisal process.

4.2 Situated Cognizers

The view of emotion as flexible adaptation mechanism is fully compatible with the concept of a situated agent. Agre [1], one of the proponents of situated AI, argues that the interaction between an agent and its environment should guide the analysis of living agents and the design of artificial ones. In accordance with this view, Lazarus ([38], p. 29) identifies the “person-environment relationship” as the “basic arena of analysis for the study of the emotion process”. Appraisal is the evaluation of this relationship, while action tendencies and coping are aimed at changing this relationship (Frijda [28] also uses the term “relational action tendencies”).

A guideline for modelling appraisal in a situated agent may be the concept of *affordances*. An affordance is defined by Gibson ([32], p. 127) as “what it offers the animal, what it provides or furnishes, either for good or ill”. The general idea is that a perceptually attuned animal actively perceives meaning in the environment without further interpretative cognitive processing. So there is a direct coupling between perception and action characteristic for situated agents. The field of “ecological psychology” is based on this idea. Both Frijda [29] and Lazarus [38] emphasize that appraisal may be based on the perception of affordances. An example of a system using a concept closely related to affordances is the PARETO plan execution system by Pryor [52]. PARETO is based on the RAP system [26] and uses “reference features” for recognizing and reasoning about opportunities.

Baron [5, 44] applies the affordance concept to social perception, e.g. to emotion perception, impression formation, and causal attribution. He establishes a “dual-mode theory of social knowing” [5], in which the ecological view and the traditional cognitive view are treated not as contradictory, but as complementary. The ecological view emphasizes the direct detection of information given in the stimulus configuration. In contrast, the cognitive view emphasizes the knowledge representations and inference processes people use to elaborate and interpret the stimulus.

We apply this dual-mode view to the nature of appraisal. Although appraisal

may be based on the direct perception of (social) affordances, it nevertheless involves *cognizing*. Chomsky [15] introduces this term to denote a person's relation to his or her knowledge. Cognizing means having access to knowledge that is not necessarily accessible to consciousness and does not consist in warranted or justified belief. A typical example of cognized knowledge is a speaker's knowledge of grammar. As described in section 2.3, the appraisal process is not necessarily conscious or voluntary, but nevertheless based on knowledge, e.g. represented in form of schemata. So, besides the direct perception of affordances, cognizing is a central part of the appraisal process.

We chose the term *situated cognizers* to emphasize the importance of both situatedness and cognizing for the view of emotion as flexible adaptation mechanism adopted in this paper.

4.3 Tractability

An interesting and important aspect regards the *tractability* of the architecture for situated cognizers, particularly with respect to the design of flexible representation schemes that are perceptually tractable in the sense exposed by Ian Horswill in [36]. Horswill specifically addresses the problem of maintenance of complex world models with limited resources, i.e., “the perception system's problem of knowing how and when to allocate its internal resources to updating different components of the world model” ([36], p.290). While Horswill envisages solutions based on mixed reaction and deliberation, he does not hesitate to express reservations with respect to simple layering approaches. Instead, he finds that “In practice, hybrid systems work well when they are given problems in which the deliberative components (1) are given little to do and (2) work only on problems, such as topological path planning, in which the relevant portions of the world model are updates slowly or not at all. . . . I believe that the solution to perceptual and combinatoric problems lies not in adjoining traditional and reactive systems but in finding useful intermediate points in between.” (*ibid.*, p.290).

It is our intention to devise TABASCO as another derivative of the 3T architecture that Horswill cites among the examples of existing “useful intermediate points”: smart and programmable reactive systems (e.g. sensory-motor mechanisms) that are allowed to spawn off graph search algorithms or resource allocators (in our case e.g. coping strategies, action tendency generators) as primitive operations when necessary. This design methodology is also in accordance with the bottom-up Animat/Artificial Life approach mentioned in section 3.2.

4.4 A Software Architecture

TABASCO is intended as an architecture for software agents situated in virtual environments. We can put forward the following points in support of our “software-only” approach:

- The discussion between proponents of software and hardware approaches (e.g., the dispute between Etzioni and Brooks) has been settled, with the conclusion that both fields offer their respective rich bag of research problems (see [35]).
- Bridging and connecting the two fields of software and hardware agent research, there is an important common stream of research concerned with software architectures that are increasingly used for hardware and software agents alike; witness the convergence on layered—often “three level”—architectures [34].
- Importantly, many potential reservations from an Artificial Life point of view can be done away with by taking care in providing an environment that on one hand is sufficiently rich in detail ([2], [45]), and on the other hand comprises software agents modeled at an appropriate level of internal (information processing) and external (bodily) complexity.
- Finally, the result of this research could always be valued as a necessary and thus worthwhile first step towards a realization of embodied solutions [30].
- In any case, it also remains an established fact that the development of hardware embodiments is substantially costlier in terms of required infrastructure, variety of required expertise, and project time.

5 Summary and Future Work

In this expository paper, we have sketched out some of the background that we are drawing upon in the attempt to design TABASCO, a tractable appraisal-based architecture for situated cognizers. We have tried to cast some light on the relevance of each of the aspects identified in the project’s name. In the best of the ALife tradition, we do expect the strong links we tied between (psychological) emotion research and work on situated autonomous agents to bring forward interesting insights and further research in both areas.

Clearly, a challenging part of the detailing of the encompassing abstract design will be the identification of a limited set of functionalities allowing to proceed

to concrete implementations putting our architecture to the test: to this end, two different—if related—variants of the architecture are being worked on, necessarily hand in hand with respective virtual environments, one text-based and one graphical/3D.

References

- [1] Agre P.E.: Computational research on interaction and agency, *Artificial Intelligence*, 72(1/2):1–52, 1995.
- [2] Agre P., Horswill I.: Lifeworld analysis, *Journal of Artificial Intelligence Research*, 6:111–145, 1997.
- [3] Anderson J.R.: *The architecture of cognition*, Harvard University Press, Cambridge, MA, 1983.
- [4] Balkenius C.: Natural intelligence in artificial creatures, *Lund University Cognitive Studies*, 37, 1995.
- [5] Baron R.M.: An ecological framework for establishing a dual-mode theory for social knowing, in Bar-Tal D. & Kruglanski A.W. (eds.), *The social psychology of knowledge*, Cambridge University Press, Cambridge, UK, 48–82, 1988.
- [6] Bartlett F.C.: *Remembering: A study in experimental and social psychology*, Cambridge University Press, New York, 1932.
- [7] Blumberg B.: Action-selection in Hamsterdam: Lessons from ethology, in *Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, Brighton, England, MIT Press, Cambridge, MA, 108–117, 1994.
- [8] Blumberg B.M.: *Old tricks, new dogs: Ethology and interactive creatures*, MIT, Cambridge, MA, PhD Thesis, 1997.
- [9] Bonasso R.P., Firby R.J., Gat E., Kortenkamp D., Miller D.P., Slack M.: Experiences with an architecture for intelligent, reactive agents, in Hexmoor H. (ed.), Special issue: Software architectures for hardware agents, *Journal of Theoretical and Experimental Artificial Intelligence*, 9(2/3):237–256, 1997.
- [10] Brooks R.A.: A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.
- [11] Brooks R.A.: Intelligence without reason, in *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, Morgan Kaufman, San Mateo, CA, 569–595, 1991.
- [12] Brooks R.A.: Intelligence without representation, *Artificial Intelligence*, 47:139–159, 1991.

- [13] Cañamero D.: Modeling motivations and emotions as a basis for intelligent behavior, in *Proceedings of the First International Conference on Autonomous Agents*, Marina del Rey, CA, ACM Press, New York, 148–155, 1997.
- [14] Chapman D., Planning for conjunctive goals, *Artificial Intelligence*, 32:333–378, 1987.
- [15] Chomsky N.: *Rules and representations*, Columbia University Press, New York, 1980.
- [16] Drescher G.L.: *Made-up minds*, MIT Press, Cambridge, MA, 1991.
- [17] Dyer M.G.: *In-depth understanding: A computer model of integrated processing for narrative comprehension*, MIT Press, 1983.
- [18] Earl C., Firby R.J.: Combined execution and monitoring for control of autonomous agents, in *Proceedings of the First International Conference on Autonomous Agents*, Marina del Rey, CA, ACM Press, New York, 88–95, 1997.
- [19] Elliott C.D.: *The Affective Reasoner: A process model of emotions in a multi-agent system*, Northwestern University, Evanston, IL, PhD Thesis, 1992.
- [20] Elliott C.: Using the Affective Reasoner to support social simulations, in Bajcsy R. (ed.), *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA, 194–201, 1993.
- [21] Elliott C.: Research problems in the use of a shallow Artificial Intelligence model of personality and emotion, in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, AAAI Press/MIT Press, Cambridge, MA, 9–15, 1994.
- [22] Elliott C., Rickel J., Lester J.C.: Integrating affective computing into animated tutoring agents, in *Proceedings of the IJCAI-97 Workshop “Animated Interface Agents: Making Them Intelligent”*, Nagoya, Japan, 1997.
- [23] Elsaesser C., Slack M.G.: Integrating deliberative planning in a robot architecture, in *Proceedings of the AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space (CIRFFSS '94)*, Houston, TX, 782–787, 1994.
- [24] Ferguson I.A.: *TouringMachines: An architecture for dynamic, rational, mobile agents*, University of Cambridge, UK, PhD Thesis, 1992.
- [25] Ferguson I.A.: On the role of BDI modelling for integrated control and coordinated behavior in autonomous agents, *Applied Artificial Intelligence, Special Issue: Intelligent Agents and Multi-Agent Systems Part 1*, 9(4):421-448, 1995.
- [26] Firby R.J.: *Adaptive execution in complex dynamic worlds*, Yale University, New Haven, CT, PhD Thesis, 1989.

- [27] Fiske S.T., Taylor S.E.: *Social cognition* (2nd edition), McGraw-Hill, New York, 1991.
- [28] Frijda N.H.: *The emotions*, Cambridge University Press, Cambridge, UK, 1986.
- [29] Frijda N.H.: The place of appraisal in emotion, in Frijda N.H. (ed.), Appraisal and beyond: The issue of cognitive determinants of emotion, *Cognition and Emotion*, 7(3/4):357–388, 1993.
- [30] Gat E.: On the role of simulation in the study of autonomous mobile robots, in *AAAI 1995 Spring Symposium, Lessons Learned from Implemented Software Architectures for Physical Agents*, 1995.
- [31] Georgeff M.P., Lansky A.L.: Reactive reasoning and planning, in *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Morgan Kaufmann, Los Altos, CA, 677–682, 1987.
- [32] Gibson J.J.: *The ecological approach to visual perception*, Houghton Mifflin, Boston, 1979.
- [33] Hayes-Roth B., van Gent R., Huber D.: Acting in character, in Trapp R. & Petta P. (eds.), *Creating personalities for synthetic actors*, Springer, Heidelberg/New York, LNAI 1195, 92–112, 1997.
- [34] Hexmoor H. (ed.): Special issue: Software architectures for hardware agents, *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2/3), 1997.
- [35] Hexmoor H., Kortenkamp D., Horswill I.: Software architectures for hardware agents, in Hexmoor H. (ed.), Special issue: Software architectures for hardware agents, *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2/3):147–156, 1997.
- [36] Horswill I.: Visual architecture and cognitive architecture, in Hexmoor H. (ed.), Special issue: Software architectures for hardware agents, *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2/3):277–292, 1997.
- [37] Jennings N.R., Sycara K., Wooldridge M.: A roadmap of agent research and development, *International Journal of Autonomous Agents and Multi-Agent Systems*, 1(1):275–306, 1998.
- [38] Lazarus, R.S.: *Emotion and adaptation*, Oxford University Press, New York, 1991.
- [39] LeDoux J.E.: *The emotional brain*, Simon & Schuster, New York, 1996.
- [40] Leventhal H.: A perceptual-motor theory of emotion, *Advances in Experimental Social Psychology*, 17:117–182, 1984.
- [41] Leventhal H., Scherer K.R.: The relationship of emotion to cognition: A functional approach to a semantic controversy, *Cognition and Emotion*, 1(1):3–28, 1987.

- [42] Loyall A.B.: *Believable agents: Building interactive personalities*, Carnegie-Mellon University, Pittsburgh, PA, PhD Thesis, 1997.
- [43] Loyall A.B., Bates J.: Real-time control of animated broad agents, in *Proceedings of the 15th Annual Conference of the Cognitive Science Society*, Boulder, CO, 1993.
- [44] McArthur L.Z., Baron R.M.: Toward an ecological theory of social perception, *Psychological Review*, 90(3):215–238, 1983.
- [45] McFarland D., Boesser T.: *Intelligent behavior in animals and robots*, A Bradford Book, MIT Press, Cambridge, MA, 1993.
- [46] O’Rorke P., Ortony A.: Explaining emotions, *Cognitive Science*, 18(2):283–323, 1994.
- [47] Ortony A., Clore G.L., Collins A.: *The cognitive structure of emotions*, Cambridge University Press, Cambridge, UK, 1988.
- [48] Pfeifer R.: Artificial Intelligence models of emotion, in Hamilton V., Bower G.H., Frijda N.H. (eds.), *Cognitive perspectives on emotion and motivation, Behavioural and Social Sciences*, 44:287–320, Kluwer Academic Publishers, 1988.
- [49] Pfeifer R.: The Fungus Eater approach to emotion: A view from artificial intelligence, *Cognitive Studies*, 1:42–57, 1994.
- [50] Pfeifer R.: Symbols, patterns, and behavior: towards a new understanding of intelligence, in *Proceedings of the Japanese Conference on Artificial Intelligence*, 1–15, 1996.
- [51] Picard R.W.: *Affective computing*, MIT Press, Cambridge, MA, 1997.
- [52] Pryor L.M.: *Opportunities and planning in an unpredictable world*, Northwestern University, Evanston, IL, PhD Thesis, 1994.
- [53] Van Reekum C.M., Scherer K.R.: Levels of processing in emotion-antecedent appraisal, in Matthews G. (ed.), *Cognitive science perspectives on personality and emotion*, Elsevier, Amsterdam, 259–300, 1997.
- [54] Reilly W.S.N.: *Believable social and emotional agents*, Carnegie Mellon University, Pittsburgh, PA, PhD Thesis, 1996.
- [55] Roseman I.J.: Cognitive determinants of emotion: A structural theory, in Shaver P. (ed.), *Review of personality and social psychology* (Vol. 5), Sage, Beverly Hills, CA, 11–36, 1984.
- [56] Rousseau D., Hayes-Roth B.: *Improvisational synthetic actors with flexible personalities*, Knowledge Systems Laboratory, Computer Science Dept., Stanford University, Stanford, CA, Report No. KSL 97–10, 1997.

- [57] Scherer, K.R.: On the nature and function of emotion: A component process approach, in Scherer K.R. & Ekman P. (eds.), *Approaches to emotion*, Erlbaum, Hillsdale, NJ, 293–318, 1984.
- [58] Scherer K.R.: Criteria for emotion-antecedent appraisal: A review, in Hamilton V., Bower G.H. & Frijda N.H. (eds.), *Cognitive perspectives on emotion and motivation*, Kluwer, Dordrecht, 89–126, 1988.
- [59] Sloman A.: *The computer revolution in philosophy: Philosophy, science and models of mind*, Harvester Press (and Humanities Press), Hassocks, Sussex, 1978.
- [60] Sloman A.: What sort of control system is able to have a personality?, in Trappl R. & Petta P. (eds.), *Creating personalities for synthetic actors*, Springer, Heidelberg/New York, LNAI 1195, 166–208, 1997.
- [61] Sloman A., Croucher M.: Why robots will have emotions, in *Proceedings of the 7th International Joint Conference on AI*, Vancouver, Canada, 1981.
- [62] Smith C.A., Ellsworth P.C.: Patterns of cognitive appraisal in emotion, *Journal of Personality and Social Psychology*, 48:813–838, 1985.
- [63] Smith C.A., Griner L.A., Kirby L.D., Scott H.S.: Toward a process model of appraisal in emotion, in Frijda N.H. (ed.), *Proceedings of the IXth Conference of the International Society for Research on Emotion, ISRE '96*, ISRE Publications, Storrs, 1996.
- [64] Smith C.A., Lazarus R.S.: Emotion and adaptation, in Pervin L.A. (ed.), *Handbook of personality: Theory and research*, Guilford, New York, 609–637, 1990.
- [65] Velasquez J., Maes P.: Cathexis: A computational model of emotions, in *Proceedings of the First International Conference on Autonomous Agents*, Marina del Rey, CA, ACM Press, New York, 518–519, 1997.
- [66] Wilkins D.E., Myers K.L., Lowrance J.D., Wesley L.P.: Planning and reacting in uncertain and dynamic environments, *Journal of Experimental and Theoretical Artificial Intelligence*, 7(1):197–227, 1995.
- [67] Wooldridge M., Jennings N.R.: Intelligent agents: Theory and practice, *The Knowledge Engineering Review*, 10(2):115–152, 1995.
- [68] Wright I.P.: *Emotional agents*, University of Birmingham, Birmingham, UK, PhD Thesis, 1997.
- [69] Yu S.T., Slack M.G., Miller D.P.: A streamlined software environment for situated skills, in *Proceedings of the AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space (CIRFFSS '94)*, Houston, TX, 233–239, 1994.
- [70] Zajonc R.B.: Feeling and thinking: Preferences need no inferences, *American Psychologist*, 2:151–176, 1980.