Abstraction Level Regulation of Cognitive Processing Through Emotion-Based Attention Mechanisms

Luís Morgado[1,2] and Graça Gaspar[2]

[1] Instituto Superior de Engenharia de Lisboa
Rua Conselheiro Emídio Navarro, 1949-014 Lisboa, Portugal
lm@isel.ipl.pt
[2] Faculdade de Ciências da Universidade de Lisboa
Universidade de Lisboa, Campo Grande, 1749-016 Lisboa, Portugal
gg@di.fc.ul.pt

Abstract. In domains where time and resources are limited, the ability to balance resource consumption according to the problem characteristics and to the required solution quality is a crucial aspect of intelligent behavior. Growing evidence indicates that emotional phenomena may play an important role in that balance. To support this view we propose an agent model where emotion and reasoning are conceived as two symbiotically integrated aspects of cognitive processing. In this paper we concretize this view by extending emotion-based regulation of cognitive activity to enable an active control of the abstraction level at which cognitive processes operate through emotion-based attention mechanisms, thus allowing a dynamical adjustment of the resources used. Experimental results are presented to illustrate the proposed approach and to evaluate its effectiveness in a scenario where reasoning under time-limited conditions in a dynamic environment is required.

1 Introduction

Since the early days of AI, the regulation of time and resources used by cognitive processes emerged as a fundamental aspect for practical intelligent behavior. Some 40 years ago, Simon [1] pointed out this problem and suggested that emotion could play a key role in that regulation. However, that suggestion remained almost unexplored until strong evidence from neurosciences started to emerge supporting the fact that emotion can effectively play a fundamental role in reasoning and decision-making (e.g. [2; 3]).

The need to control the cognitive processes has been mainly addressed by resorting to meta-cognition (e.g. [4]). While meta-cognition evidences relevant aspects to address the regulation of cognitive activity, in some way the problem of bounded resources is just transferred to another level of cognition, since meta-cognitive processes themselves can be computationally intensive [5]. Based on experimental evidence from neurosciences (e.g. [2; 3]) emotion-based mechanisms can constitute an interesting alternative.

Two main approaches have characterized the development of emotion models for intelligent agents, physiologically inspired models (e.g. [6; 7]) and appraisal theories inspired models (e.g. [8; 9; 10]). These models have allowed the definition of emotion-like characteristics and behavior, however they also have some drawbacks. Physiologically inspired models are based on specific mechanisms of biological organisms, like hormonal mechanisms, resulting in highly specific implementations limited to relatively simple agents and contexts. According to appraisal models, emotional phenomena derive from appraisal processes based on specific appraisal dimensions, which leads to a view of emotion characterized by discrete emotional qualities.

A distinctive aspect of our approach is the fact that emotional phenomena are modeled in a way that preserves their dynamic and continuous nature and their double role as both a contributing factor and a result of cognitive activity, therefore allowing a tight integration of emotion and cognition and enabling an adaptive regulation of the cognitive processes of an agent. At the core of that regulation are emotion-based attention mechanisms, which dynamically focus cognitive processing. To concretize this approach we developed the *flow model of emotion* and the *agent flow model*, which will be briefly presented in sections 2 and 3.

Our previous work [20; 28] had already addressed the adaptive regulation of cognitive processes by defining two mechanisms: an attention field mechanism that filters the cognitive elements over which processing can occur, and a temporal focusing mechanism that regulates the time available for cognitive processing. In this paper we refine the focusing mechanisms of the agent flow model and extend them in order to adaptively control the abstraction level at which the agent operates.

The variation of the abstraction level at which the world is modeled is a powerful technique to reduce resource use and thus increase scalability, whose results have been demonstrated in areas such as control learning [11], reinforcement learning [12] and search algorithms [13]. We show that the emotion-based attention mechanisms can also be used effectively to adjust the abstraction level adopted by the agent, in unknown dynamic environments, and can be incorporated with other emotionally guided focusing mechanisms in order to increase performance.

In section 5, we will report experimental results that illustrate the use of those mechanisms in the Tileworld scenario and in section 6 we discuss related work and draw some conclusions and directions for future research.

2 The Flow Model of Emotion

The assumption that emotions can be divided into discrete and independent categories has encouraged the structural characterization of emotion based on linguistic labels. That has been the main approach to model emotion for the design of artificial agents. However the limitations of that approach are increasingly recognized, especially in what refers to modeling the dynamic non-linear aspects of emotional phenomena.

Alternative views have been proposed (e.g. [14, 15]) that consider those dynamic aspects. However they maintain a typical commitment to an anthropomorphic view of emotion, which leads to complexity, brittleness and lack of flexibility in agent design and implementation, especially when we want to model agents of different kinds and levels of complexity. Due to the prevalence of this anthropomorphic view, some aspects of emotional phenomena have been largely overlooked that could provide interesting directions to address these issues.

One of these aspects is the evidence that affective-emotional phenomena are pervasive among biological organisms, even the simplest ones such as unicellular organisms (e.g. [29]). Taking this observation as a starting point, affective-emotional phenomena are not necessarily dependent on cognitive or even nervous structures, but on more basic biophysical principles.

Following this line of research we adopt a view where "basic biological organization is brought about by a complex web of energy flows" [30]. To support this view an agent is modeled as a *dissipative structure* [31]. Dissipative structures are open systems governed by the interchange of energy with the environment and able to maintain themselves in a state far from equilibrium, yet keeping an internally stable overall structure. The maintenance of that internal stability in spite of environmental changes is done through feedback networks that motivate the system to act. The maintenance of a basic life support energy flow can be seen as a base motivation.

Even though motivations can take various forms according to the cognitive context (e.g. drives, desires), in any case, to achieve its motivations an agent must be able to produce the adequate change in the environment, by applying an internal potential. However, the concretization of the intended change depends on the characteristics of the current environmental situation that, from a thermodynamic point of view, can be modeled as an agent-environment coupling conductance. Therefore, the agent-environment relation can be modeled as a relation between an agent's internal potential, its *achievement potential*, and the agent-environment coupling conductance, the *achievement conductance*. The achievement potential represents the potential of change that the agent is able to produce in the environment to achieve the intended state-of-affairs. The achievement conductance represents the degree of the environment's conduciveness or resistance to that change.

From a thermodynamic point of view, the achievement potential can be viewed as a force (P) and the achievement conductance as a transport property (C). The behavioral dynamics of an agent can therefore be characterized as a relation corresponding to a flow, called *achievement flow* (F), which results from the application of a potential P over a conductance C. The behavioral forces that arise from this dynamic relation between achievement potential and achievement conductance, expressed as energy flows, generate behavioral dynamics that underlie the cognitive activity of an agent. In the proposed model we consider emotional phenomena as the expression of those dynamics.

Those dynamics are described by a vectorial function ED, called *emotional disposition*, defined as:

$$ED \equiv (\delta P, \delta F) \quad \text{where} \quad \delta P = \frac{dP}{dt} \quad \text{and} \quad \delta F = \frac{dF}{dt} \tag{1}$$

This notion of *emotional disposition* is defined as an action regulatory disposition or tendency, but it does not constitute in itself an emotion. In the proposed model, phenomena such as emotions and moods arise from an interdependent relation between emotional dispositions and cognitive activity, according to the agent-environment interaction patterns and the cognitive context (e.g. self-reflective or social context).

3.1 Emotional Qualitative Characterization

A basic problem of emotion modeling is how to explain in a single model both the dynamic, continuously fluctuating nature of emotion processes and the existence of discrete labels referring to steady states [14]. To address these questions we need to make a qualitative characterization of emotional dispositions.

As can be seen in figure 1.a, at a given instant $t = \tau$ an emotional disposition vector has a quality, defined by its orientation (or argument) and an intensity defined by its module. Each quadrant of the two dimensional space $\delta P \times \delta F$ can be directly related to a specific kind of *emotional disposition quality* [33] as indicated in figure 1.b. As an example, quadrant Q–III ($\delta P < 0$ and $\delta F < 0$) corresponds to situations where the agent does not have the capacity to handle the “adversities”, which are typically fear situations.

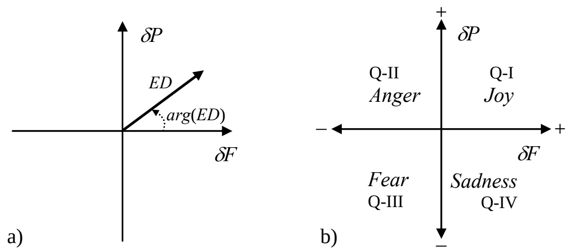

Fig. 1. Vector ED as a function of δP and δF (a); relation between ED quadrants and emotional quality tendency (b).

It is important to note that the emotional tendency associated to each quadrant (joy, anger, fear, sadness) is only indicative of its main nature, since the quality of the emotional disposition is continuous. This is consistent with phenomenological well-known emotion blends.

4 The Agent Flow Model

Although inspired by biophysical analogies, the main aim of the proposed model is to support the development and implementation of artificial agents, independently of their kind or level of complexity. Therefore it is necessary to concretize the base notions of the model in a computationally tractable way.

The first aspect that we need to address is the notion of energy. In thermodynamics, energy is usually defined as the capacity to produce work. In the context of the flow model, energy can be defined as the capacity of an agent to act or, in a wide sense, to produce change. Considering an agent as a dissipative structure, that change is oriented towards the achievement of motivations driven by internal potentials and expressed through energy flows. That is, both the agent and the environment can be modeled as a composition of multiple energetic potentials with different characteristics. In this way, the notion of energetic potential is the base notion that allows unifying the different aspects that characterize an agent in a single uniform framework.

4.1 Agent Cognitive Structure

Energetic potentials can aggregate to form composite potentials. These aggregated potentials can represent different elements of an agent, such as a perception, a memory or an intention, providing an adequate support to model the cognitive structure of an agent. Therefore they are generically called *cognitive elements*. The potentials that form cognitive elements express aspects that the agent is able to discriminate and perceive, such as "weight" or "color", commonly called quality dimensions [18].

Formally, cognitive potentials are modeled as a composition of two types of signals: a base signal $\varphi(t)$ with a specific angular frequency ω that identifies the discriminated aspect or quality; and a quantitative signal $\rho(t)$ corresponding to the actual value of the discriminated quality, expressed as a frequency shift $\Delta\omega$ that modulates the base signal $\varphi(t)$. That is:

$$p(t) = \rho(t).\varphi(t) \tag{2}$$

Through superposition, aggregates of potentials can be formed. Superposition is possible because the base signals that characterize the cognitive potentials are orthogonal among each other, which implies superposition of energy. Therefore a cognitive element $\sigma(t)$ is defined as a superposition of cognitive potentials. That is:

$$\sigma(t) = \sum_{i=1}^{K} p_i(t) \tag{3}$$

where K is the number of potentials in the aggregate.

Cognitive elements play different roles in cognitive activity. Three main roles can be identified: *observations*, *motivators*, and *mediators*. *Observations* result from perception processes, representing the current environmental situation. They can also result from simulated experience [28]. *Motivators* and *mediators* are formed internally or embedded in agents' structure. *Motivators* represent intended situations, acting as motivating forces driving agent's behavior. *Mediators* describe the media that supports action, forming an interface between internal cognitive processing and action. For instance, planning processes produce sequences of *mediators* that are translated by action processes into concrete action.

4.2 Cognitive Space

To describe the structural and dynamic aspects of the proposed model in a concise way we can observe that the base signals that compose potentials and cognitive elements form a signal space underlying the cognitive structure of an agent, which we call a *cognitive space*. Formally, a cognitive space CS_K is defined by a set of K orthonormal basis functions $\Phi = \{\varphi_i: i = 1, 2, \ldots, K\}$ with $K \in \mathbb{N}$. Each basis function φ_i corresponds to a base signal $\varphi_i(t)$ with a specific quality ω_i.

Cognitive elements correspond to specific positions in the cognitive space. Since cognitive elements change with time, at successive time instants they occupy different positions, describing trajectories that reflect the behavior of an agent. At some instant $t = \tau$, a cognitive element $\sigma(t)$ is represented in a cognitive space CS_K as a vector $\boldsymbol{\sigma}$, defined as:

$$\boldsymbol{\sigma} = (\rho_0, \rho_1, \ldots, \rho_k) \quad \mathbf{(4)}$$

where the dimensional factors $\rho_i \in \mathbb{C}$ convey the intensity and frequency shift of quality ω_i in the cognitive element. Besides enabling a concise description of agents' cognitive structure, the cognitive space also enables a concise description of cognitive dynamics as movement of cognitive elements, as will be discussed next.

4.3 Cognitive Dynamics

The interactions between cognitive elements act as behavioral driving forces, forming the basis of cognitive dynamics.

One of the main characteristics of intelligent behavior is the orientation towards the achievement of motivations. This process of motivation achievement can be described through the evolution of the relation between the current situation, represented by an *observation*, and an intended situation, represented by a *motivator*. The cognitive activity of an agent is consequently guided by the maximization of the flows that lead to the reduction of the distance between observations and motivators, through the use of *mediators*. This process can be pictured in the cognitive space, where motivators and observations correspond to specific positions and mediators define directions of movement, as illustrated in figure 2.

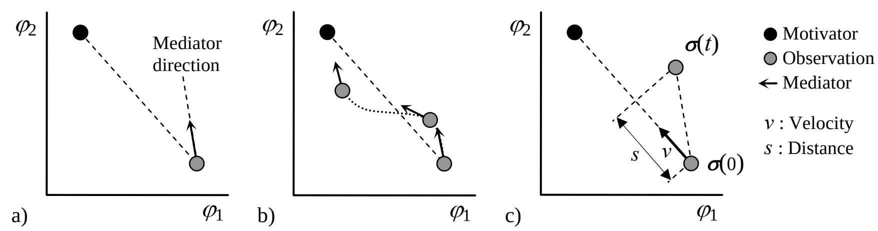

Fig. 2. Elements participating in the achievement of a motivator in a two-dimensional cognitive space.

Figure 2.b shows a possible trajectory resulting from the adjustment of agent's behavior to changes in the environment, by switching to different mediators. Independently of the specific processes that generated the new mediators, from the point of view of the proposed model emotional phenomena are considered the expression of the forces that led to that change, characterized as emotional dispositions.

In the cognitive space, the cognitive dynamics can be described by the movements of cognitive elements, and the associated emotional dispositions can be defined by the evolution of the distance covered s and the velocity v of the observation, relative to the motivator (figure 2.c). That is:

$$ED \equiv (\delta s, \delta v) \quad \text{where} \quad \delta s = \frac{ds}{dt} \quad \text{and} \quad \delta v = \frac{dv}{dt} \tag{5}$$

These emotional dispositions represent behavioral forces that constrain the cognitive processes of an agent. Therefore, they are, at the same time, a result of the cognitive activity and a constraint that influences it, reflecting the symbiotic relation between emotion and cognition.

3.3 Emotional Disposition Mechanisms

Emotional disposition mechanisms detect the emotional disposition dynamics, previously described, producing concrete signals. Given an observation and a motivator, two affective signals are generated, λ^+ and λ^-, that convey the affective character underlying those emotional dispositions.

The emotional disposition components δs and δv (5) convey a hedonic quality, associated to the increase or decrease of agent's well being in relation to the achievement of its motivators. An increase of s that is larger than the decrease of v ($\delta s > -\delta v$) expresses an improvement in the achievement conditions, corresponding to a positive valence. The opposite relation means a deterioration of the achievement conditions, corresponding to a negative valence. This valence aspect, pleasant vs. unpleasant or positive vs. negative, constitutes an *affective* quality [2]. The most favorable affective situation occurs when both δs and δv have positive values. In the emotional disposition plane this situation corresponds to ED vectors located in quadrant Q-I. In the same way, the most unfavorable affective situation occurs when both δs and δv have negative values, corresponding to ED vectors located in quadrant Q-III.

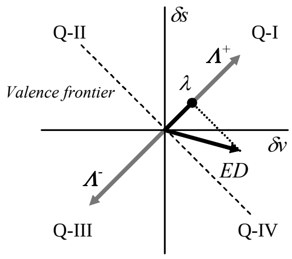

Fig. 3. Affective characterization in the emotional disposition plane.

Considering normalized valence values, a positive valence is represented by the projection of the ED over the reference vector $\boldsymbol{\Lambda}^+ = (1,1)$ and a negative valence by the projection over $\boldsymbol{\Lambda}^- = (-1,-1)$, as shown in figure 3. That is:

$$\lambda^+ = \text{proj}(ED, \boldsymbol{\Lambda}^+) \quad \text{and} \quad \lambda^- = \text{proj}(ED, \boldsymbol{\Lambda}^-) \tag{6}$$

This geometric characterization corresponds in practice to a linear composition of the δs and δv components.

4 Adaptive Regulation of Cognitive Activity

In real-world domains resources are limited, change is pervasive and behavior must be produced in real-time to cope with that change. This leads to a tradeoff between resource consumption and solution quality that has been commonly associated to reasoning and decision-making processes. However, it can involve the overall cognitive activity, for instance, action can be more or less precise, perception more or less encompassing, memory formation more or less detailed. On the other hand, this tradeoff must reflect the dynamic relation between the agent and the environment and therefore should be adaptively regulated. Emotional phenomena is used in the agent flow model to support that adaptive regulation.

The regulation of the cognitive activity according to the present achievement conditions involves focusing cognitive processes along two perspectives: (i) a *spatial perspective* that refers to the space of cognitive elements over which processing can occur; and (ii) a *temporal perspective* that refers to the time available for cognitive processing. These two perspectives are concretized by two main mechanisms, an *attention focusing* mechanism and a *temporal focusing* mechanism. Figure 4 illustrates how these mechanisms are interrelated.

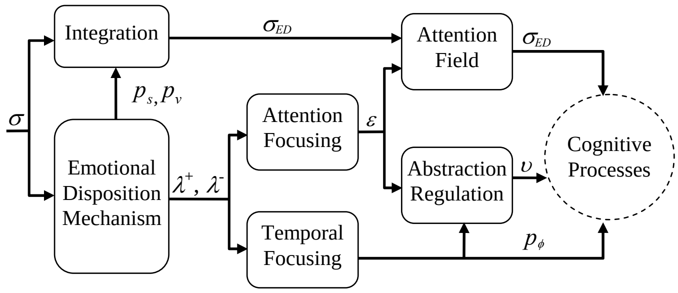

Fig. 4. Mechanisms underlying the cognitive activity regulation.

Given an observation and a motivator, the emotional disposition mechanism generates two types of signals: emotional disposition cognitive potentials, p_s and p_v, which convey the emotional disposition; and affective signals, λ^+ and λ^-, that convey the affective character underlying those cognitive potentials. These regulatory signals λ^+ and λ^- are directly input to both focusing mechanisms. On the other hand, the p_s and p_v cognitive potentials get integrated with the pair of cognitive elements that originated them, an observation and a motivator, constituting a composite cognitive element $\boldsymbol{\sigma}_{ED}$ with an emotional disposition content.

Attention focusing and temporal focusing constitute a primary level of regulatory mechanisms, that are structurally similar. The attention field and the abstraction regulation mechanism constitute a second level of regulation feeding on either or both of the previous.

4.1 Temporal Focusing

The temporal focusing mechanism controls the rate of cognitive activity by providing a time-base for overall cognitive processing, corresponding to a signal p_ϕ with frequency ω_ϕ. That is, it determines the time available before some behavior must be produced. There is experimental evidence that emotions affect the perception of time and also that they affect the pressure to cope with situations (e.g. [32]). Based on these evidences, we suggest that emotions may lead to an adjustment in the rate of cognitive activity. To capture this influence, we consider that the frequency ω_ϕ expresses the cumulative effect of λ^+ and λ^- signals, reflecting the prevailing affective character of the achievement conditions. We consider the cumulative effect in order to model more stable emotional patterns instead of the instantaneous, continuously changing, emotional tendencies resulting from isolated experiences. That is:

$$\frac{d\omega_\phi}{dt} = \beta^+.\lambda^+ + \beta^-.\lambda^- \tag{7}$$

where the sensitivity coefficients $\beta^+ \in \mathbb{R}$ and $\beta^- \in \mathbb{R}$ determine the influence of λ^+ and λ^- signals, respectively.

Temporal focusing allows taking advantage of different types of bounded reasoning mechanisms, such as partial planning (e.g. [20]).

4.2 Attention Focusing

Some theories of perception [21] have identified three main aspects of attention: (i) the *locus*, that is, the cognitive elements at which cognitive activity is directed; (ii) the *extent*, that is, the range of cognitive elements involved; and (iii) the *detail level* of those cognitive elements.

In the proposed model, both the extent of attention focusing and the detail level at which cognitive processing will occur depend on a signal ε, produced by the attention focusing mechanism. This signal results from the cumulative effect of λ^+ and λ^- signals, in a similar way to the temporal focusing signal. That is:

$$\frac{d\varepsilon}{dt} = \alpha^+.\lambda^+ + \alpha^-.\lambda^- \tag{8}$$

where the sensitivity coefficients $\alpha^+ \in \mathbb{R}$ and $\alpha^- \in \mathbb{R}$ determine the influence of λ^+ and λ^- signals, respectively.

4.3 Attention Field

The attention field mechanism acts like a depletion barrier, producing an *attention field* formed by the cognitive elements able to bypass the barrier. Only the elements in the attention field are considered by the high-level cognitive processes, such as reasoning and deliberation. The depletion barrier is characterized by an intensity, given by the ε signal above described, and a *permeability* μ. This permeability determines the intensity ε^σ of the interaction between a cognitive element σ and the depletion barrier, defined as:

$$\varepsilon^\sigma = \mu_s.p_s^\sigma + \mu_v.p_v^\sigma \tag{9}$$

where $\mu_s \in \mathbb{R}$ and $\mu_v \in \mathbb{R}$ are coefficients that determine the influence of p_s and p_v cognitive potentials. Given a certain depletion intensity ε, a cognitive element σ bypasses the barrier and is included in the attention field if $\varepsilon^\sigma > \varepsilon$.

4.4 Abstraction Level Regulation

Abstraction plays a key role to handle computational complexity in the design of AI systems. Different approaches to abstraction have been adopted, such as constraint relaxation, where a problem is simplified by eliminating some conditions, as is the case in hierarchical planning (e.g. [22]), or action abstraction, as used for instance in reinforcement learning, where multiple actions are abstracted into one multi-step action (e.g. [12]). However, this kind of abstraction techniques needs to be regulated in order to be effective. For instance, an inadequate level of detail could lead to aliasing phenomena where relevant details of the processed representations are overlooked.

To regulate the abstraction level of cognitive processing we need to consider the integrated operation of both attention and temporal focusing mechanisms. Together, they act like a zoom control, leading the cognitive processes to increase or decrease the detail level. On one side, an increase in the rate of cognitive activity ω_ϕ results in less time available for cognitive processing, prompting for increased abstraction to reduce the computational effort. On the other side, an increase in the attention focusing signal ε prompts for increased detail. The abstraction level υ that results from these two complementary influences is defined as follows:

$$\upsilon = \omega'_\phi.(1-\varepsilon') \qquad \textbf{(10)}$$

where ω'_ϕ and ε' are the normalized values of ω_ϕ and ε, respectively, in the range [0,1]. In this way, the potentially conflicting influence of both these mechanisms is balanced.

5 Experimental Results

To illustrate the operation of the above described mechanisms, we will consider a scenario where reasoning under time-limited conditions in a dynamic environment is required. The experimental framework is an implementation of the *Tileworld* domain that follows the specification presented in [23], also adopted by Schut *et al.* [24], which will provide reference points for comparison of results.

The Tileworld is a 2-dimensional grid on which an agent scores points by moving to targets, known as holes. When the agent reaches a hole, the hole is filled and disappears. Holes appear in specific instants, in randomly selected empty squares, and exist for a length of time. Both holes' gestation time and life expectancy are taken from given independent random distributions unknown to the agent. The task of the agent is to visit holes in order to score as many points as possible.

In our implementation, each hole perceived by the agent is modeled as a motivator, the current position of the agent is modeled as an observation, and the possible actions are modeled as mediators. As in Kinny and Georgeff's implementation, agents only generate plans for visiting a single hole, rather than planning multiple-hole tours.

The planning process is based upon a state space A^* planner. To enable the interruption of the planning process due to temporal focusing, partial planning is supported. To enable the regulation of the abstraction level at which plans are formed, multi-step actions are supported. A set of multi-step actions of degree $n \in \mathbb{N}$ is defined as $A^{(n)} = \{a^n \mid a \in A^{(1)}\}$, where a^n denotes the multi-step action that results if action a is executed n consecutive times and $A^{(1)}$ denotes the set of primitive actions [12]. The degree n is determined by the abstraction level υ multiplied by a scaling factor. The planner is complete, in the sense that if no solution is found at some action degree n it will decrease the action degree until a solution is found or no solution is possible, although optimal solution is not guaranteed.

The attention field constrains the set of motivators (holes) over which deliberation will occur, guiding the reconsideration of intentions, that is, the choice of the motivator that the agent will try to reach (the current intention). It also controls the switch between planning and action activities. While the motivator corresponding to the current intention remains in the attention field, no reconsideration will occur and action is activated according to the previously determined plan. Otherwise, the motivator closest to the agent position is selected as the next current intention and planning is activated.

The switch between planning and acting is also determined by the temporal focusing mechanism. If during the activation of planning a period of cognitive activity ends, planning is interrupted, the best partial plan found so far is considered, and the first planned action is executed.

5.1 Results and Analysis

A set of experiments will be reported concerning the operation of the attention focusing and temporal focusing mechanisms to regulate the abstraction level at which the planning process operates. The permeability and sensitivity coefficients of these mechanisms were considered fixed parameters of the implementation. Two dependent variables were measured: (i) the effectiveness of the agent, defined as the ratio of the actual score achieved by the agent to the maximum score that could in principle have been achieved; and (ii) the total planning cost, defined as the sum of the planning costs for all plans generated during a run. For a given plan, the planning cost is proportional to the dimension of the search tree explored during plan formation. In each experiment the *dynamism* of the environment (γ), denoted by an integer in the range 1 to 100, representing the ratio between the world clock rate and the agent clock rate, was varied. The results presented for a given dynamism value are the average over 100 runs of 20000 time-steps per run.

To provide a reference point for comparison of the results produced with our approach, we also implemented a purely deliberative agent following the Schut, Wooldridge and Parsons's (SWP) best reconsideration policy (from [24]). This policy lets the agent deliberate when a hole appears that is closer than the intended hole (but not on the path to the intended hole), and when the intended hole disappears. This policy improves the best policy presented by Kinny and Georgeff [23].

We will present two sets of experiments. In the first set, we compared the SWP agent with versions of our agent without using the abstraction regulation mechanism. In the second set, we compared the SWP agent and the best configuration found in the first experiments with our agent using the abstraction regulation mechanism.

Figure 5 shows the results obtained from the different experiments concerning 4 agents: the SWP agent and 3 agents corresponding to different configurations of our agent architecture. The following configurations were tested: ATF agent, with both focusing mechanisms enabled; AF agent with only the attention focusing mechanism enabled; and TF agent, with only the temporal focusing mechanism enabled.

As can be observed in figure 5, the effectiveness results for the AF agent are better than the results of the reference SWP agent, particularly for low to medium values of dynamism. For the TF agent the results are also better, except for low values of dynamism ($\log10(\gamma) < 0.8$). In the case of the ATF agent (both attention focusing and temporal focusing enabled) there is consistent improvement through the whole range of dynamism values.

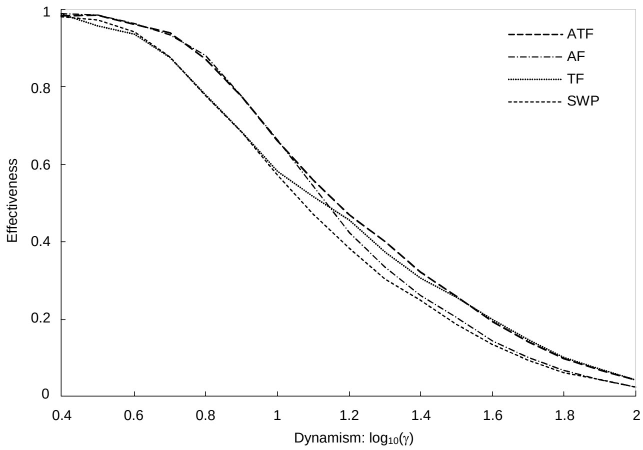

Fig. 5. Effectiveness results for different degrees of dynamism of the environment.

These results can be explained by the combined operation of emotional disposition and focusing mechanisms. The emotional disposition mechanisms provide the base support for deciding when to change the current intention by modulating the cognitive elements relevance. For instance, when a new hole appears a strong favorable emotional disposition is produced enabling the corresponding motivator to enter the attention field. In addition, the affective signals that are generated regulate the attention and temporal focusing. For instance, when the dynamism is low to medium, an agent has time to fill most of the holes that appear, and so a positive affective character prevails, which results in a low depletion intensity. Therefore, most of the motivators are present in the attention field, including the current intention. In this way, while the SWP agent reconsiders every time a hole appears that is closer than the intended hole, except if it is on the path to the intended hole, our AF and ATF agents will rarely reconsider. This reduces planning time consumption (planning cost) increasing the time available for acting. The overall result is an improved effectiveness.

In the second set of experiments, we considered two different configurations of the agent flow model architecture: a standard ATF configuration, with both focusing mechanisms enabled, as presented above, and an ATF-A configuration, which combines the attention and temporal focusing with the abstraction level regulation, as discussed in the previous section. Figures 6 and 7 show the results obtained for these three agent configurations.

As can be observed in figure 6, the effectiveness results for the ATF-A agent are considerably better than the results of the reference agents SWP and ATF in the medium range of dynamism ($0.6 < \log_{10}(\gamma) < 1.4$). The planning cost also shows improvement over the same range.

When the dynamism is low, an agent has time to fill most of the holes that appear, and so a positive affective character prevails, which results in low ε and low ω_ϕ. When the dynamism increases, agents start failing to reach some holes, and so a negative affective character raises up, which results in the increase of both ε and ω_ϕ. The increase of ω_ϕ results in shorter cognitive processing periods, restricting planning time and increasing the responsiveness of the agent to the fast changing conditions, leading to an increase in the effectiveness.

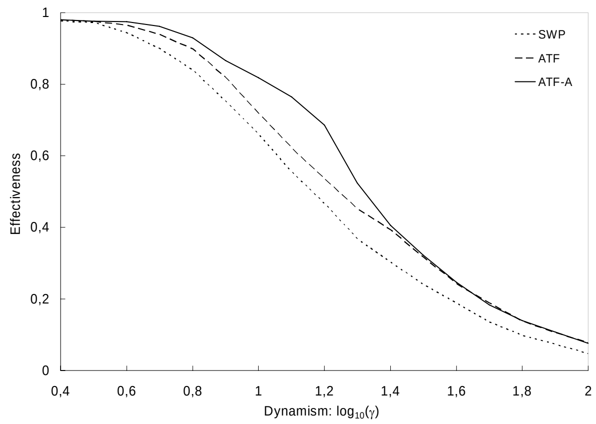

Fig. 6. Effectiveness results for different degrees of dynamism of the environment.

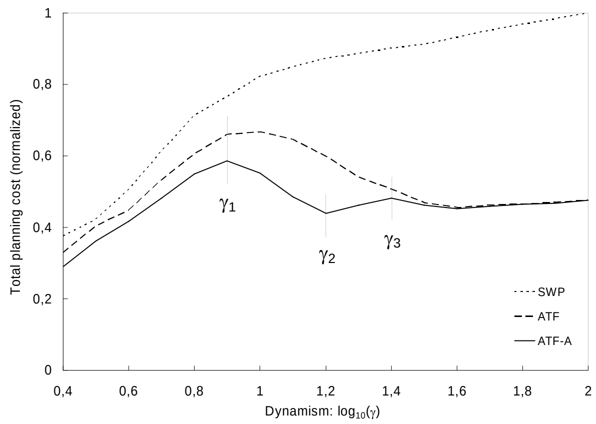

Fig. 7. Total planning cost for different degrees of dynamism of the environment.

In the case of ATF-A agent, in the medium range of dynamism, particularly after γ_1, the influence of ω'_ϕ over the abstraction level dominates that of ε', leading to an increase in the average action degree used by the planner. This reduces planning time consumption (planning cost) by forming plans with less detail, as shown in figure 7, improving its effectiveness in relation to the ATF agent. However, as the dynamism increases and more holes appear near the agent, the need for plan refinement at lower abstraction levels increases, leading to an increase of the planning cost. This is clearly noticeable after γ_2. After that point, ε' starts dominating the influence of ω'_ϕ, leading to a reduction of the abstraction level before planning starts, therefore avoiding the need for plan refinement during planning. This is noticeable after γ_3, where the ATF-A agent results converge to the ATF agent results.

In this way, the integration of attention and temporal focusing for abstraction level regulation is an effective complement to the base regulation mechanisms, providing an adequate support to control the reasoning processes.

6 Discussion and Conclusions

The main line of research on emotion for AI systems has focused on a discrete structural characterization of emotional phenomena, managing nevertheless to support the simulation of some aspects of emotional phenomena, namely the relation with cognitive functions (e.g. [10]) or the design of agents that convey a sense of emotion (e.g. [9]). However, a dynamic characterization of the relation between emotion and cognition can be particularly relevant for the design of agents able to control the reasoning processes to cope with limited time and resources. In fact, this could be a key role of emotional phenomena even in biological systems, as Simon [1] pointed out.

Our proposal departs from the main approaches to emotion modeling by considering emotion and cognition as two symbiotically integrated aspects of agent cognitive activity. This means that the relation between emotion and cognition occurs not only at a functional specialization level. Instead it is intrinsic to all cognitive activity and to the nature of the involved cognitive elements. Recent experimental results support this view, indicating that in humans, emotion and higher cognition can be truly integrated, that is, at some point of processing, functional specialization is lost and emotional and cognitive influences inseparable [3].

In this paper, we have shown how to improve the results obtained with the regulatory mechanisms of the agent flow model, a model that builds on that dynamic view. At the core of that regulation are emotion-based attention mechanisms, which focus cognitive processing. This relation between emotion, attention and cognition has been increasingly evidenced by experimental results from neurosciences (e.g. [2]).

The results presented concern the use of that abstraction regulation to control the detail of planning steps. However, this same mechanism could be used, for instance, to control the detail of perception or the granularity of memory formation. The effect of emotion-based abstraction regulation over cognitive processes other than reasoning is a subject for future work. The effect of attention based mechanisms in social interaction, as a capacity to “tune in” to others and perform joint attention behaviors [Ziemke] is also a subject for future work.

In what concerns reasoning processes, this kind of regulation is not necessarily opposed to other approaches to real-time bounded reasoning. Namely, it can feed planning based on real-time search algorithms (e.g. [25]) instead of A^*, or be integrated with planning using map abstraction (e.g. [13]) or other techniques for meta-level control of anytime algorithms (e.g. [26]). This is another area to explore in the future.

References

1. Simon, H.: Motivational and Emotional Controls of Cognition, Psychological Review, 74, 1967.
2. Damásio, A.: A Second Chance for Emotion, Cognitive Neuroscience of Emotion, eds. R. Lane, L. Nadel, Oxford Univ. Press, 2000.
3. Gray, J., Braver, T., Raichle, M.: Integration of Emotion and Cognition in the Lateral Prefrontal Cortex, Proceedings of the National Academy of Sciences USA, 2002.
4. Anderson, M., Oates, T. (eds.): Metacognition in Computation, AAAI Spring Symposium, Technical Report SS-05-04, AAAI Press, 2005.
5. Gigerenzer, G., Selten, R. (eds.): Bounded Rationality: The Adaptive Toolbox, MIT Press, 1999.
6. Cañamero, L.: Designing Emotions for Activity Selection in Autonomous Agents. In R. Trappl, ed., Emotions in Humans and Artifacts, MIT Press, 2000.
7. Almeida, L., Silva, B., Bazzan, A.: 'Towards a physiological model of emotions: first steps', AAAI Spring Symposium, ed. E. Hudlicka, L. Cañamero, Technical Report SS-04-02, 2004.
8. Ortony, A., Clore, G., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press, 1988.
9. Gratch, J., Marsella, S.: Evaluating a General Model of Emotional Appraisal and Coping, AAAI Spring Symposium, ed. E. Hudlicka, L. Cañamero, Tech. Rep. SS-04-02, 2004.
10. Hudlicka, E.: Modeling Interaction between Metacognition and Emotion in a Cognitive Architecture, AAAI Spring Symposium, Technical Report SS-05-04, AAAI Press, 2005.
11. Munos, R., Moore, A.: Variable Resolution Discretization in Optimal Control, Machine Learning,1, 1-31, 2001).
12. Schoknecht, R., Riedmiller, M.: Learning to Control at Multiple Time Scales. ICANN, 2003.
13. Sturtevant, N., Buro, M.: Partial Pathfinding Using Map Abstraction and Refinement, Proc. of the International Joint Conference on Artificial Intelligence, AAAI Press, 2005.
14. Scherer, K.: Emotions as Episodes of Subsystem Synchronization Driven by Nonlinear Appraisal Processes. In Emotion, Development, and Self-Organization, eds. M. Lewis, I. Granic, Cambridge Univ. Press, 2000.
15. Carver, C., Scheier, M.: Control Processes and Self-organization as Complementary Principles Underlying Behavior, Pers. and Social Psych. Review, 2002.
16. Kondepudi, D., Prigogine, I.: Modern Thermodynamics: From Heat Engines to Dissipative Structures, J. Wiley & Sons, 1998.
17. Morgado, L., Gaspar, G.: Emotion in Intelligent Virtual Agents: The Flow Model of Emotion. In Intelligent Virtual Agents, eds. T. Rist et al., LNAI 2792, Springer, 2003.
18. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought, MIT Press, 2000.
19. Moore, S., Oaksford, M. (eds.): Emotional Cognition, John Benjamins Press, 2002.

20. Morgado, L., Gaspar, G.: Emotion Based Adaptive Reasoning for Resource Bounded Agents, Proc. 4th Int. Joint Conference on Autonomous Agents and Multi-Agent Systems, ACM Press, 2005.
21. Coren, S., Ward, L., Enns, J.: Sensation and Perception, Wiley, 2004.
22. Ghallab, M., Nau, D., Traverso, P.: Automated Planning, Morgan Kaufmann, 2004.
23. Kinny, D., Georgeff, M.: Commitment and Effectiveness of Situated Agents, Proc. of the 12th International Joint Conference on Artificial Intelligence, 1991.
24. Schut, M., Wooldridge, M., Parsons, S.: The Theory and Practice of Intention Reconsideration. In J. Expt. Theor. Artificial Intelligence, 16:4, 2004.
25. Koenig, S., Likhachev, M.: RealTime Adaptive A*, Proc. 4th Int. Joint Conference on Autonomous Agents and Multi-Agent Systems, 2006.
26. Horvitz, E., Zilberstein, S.: Computational Tradeoffs under Bounded Resources, Artificial Intelligence Journal, 126, 2001.
27. Lindblom, J., Ziemke, T.: The Social Body in Motion: Cognitive Development in Infants and Androids, Connection Science, 18:4, 2006.
28. Morgado, L., Gaspar, G.: Adaptation and Decision-Making Driven by Emotional Memories, in Progress in Artificial Intelligence, C. Bento et al. (eds.), LNAI 3808, Springer-Verlag, 2005.
29. Staddon, J., Adaptive Dynamics: The Theoretical Analysis of Behavior, MIT Press, 2001.
30. Bergareche, A. , Ruiz-Mirazo, K.: Metabolism and the Problem of its Universalization, Biosystems, 49, 1999.
31. Nicolis, G. , Prigogine, I. : Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order trough Fluctuations. John Wiley & Sons, 1977.
32. Angrilli, A., Cherubini, P., Pavese, A.: 'The Influence of affective Factors on Time Perception', Perception & Psychophysics, 59(6), 1997.
33. Morgado, L., Gaspar, G.: Emotion in Intelligent Virtual Agents: The Flow Model of Emotion', in Intelligent Virtual Agents, eds. T. Rist et al., LNAI 2792, 31-38, Springer-Verlag, 2003.