Modeling Expectations for Affective Agents*

Michele Piunti[1], João Gonçalves[2], and Carlos Martinho[2]

[1] Institute for Cognitive Sciences and Technologies - CNR, Roma, Italy
Università di Bologna - DEIS, Bologna, Italy
michele.piunti@istc.cnr.it
[2] Instituto Superior Técnico, Universidade Técnica de Lisboa, Lisbon, Portugal
jcpg@mega.ist.utl.pt, carlos.martinho@tagus.ist.utl.pt

Abstract. In this work we discuss an enhanced architecture for cognitive agents guiding and orienting reasoning processes according to different kinds of expectations and affective states. They are assumed to involve cognitive processes including deliberation and practical reasoning where emotional biases are assumed to affect decisions, as well as situated adaptation to contexts and environment dynamism, where external events are appraised as opportunities or determent and recruits additional resources to pro-actively respond. We further provide experiments in a dynamic scenario, showing improvements in agent pro-activeness, opportunism and anticipatory effectiveness.

1 Introduction

To make a computational agent able to process representations of future states is an important challenge for forthcoming cognitive systems. In order to enhance agent's anticipatory competencies at a design level, here we refer to some of the several functions that expectations can embed in a goal directed system. The most immediate contribution of expectations is for anticipation [1]. We define *anticipation* as a behavior that does not only depend on the past and the present, but also on some knowledge about the future. Thus, anticipation is based on building internal representations of the *not yet* existent and on the ability to act, in the present, coordinating behavior and decisions with those events expected to happen in the future. Accordingly, expectations have been placed at the basis of many emotions and play a fundamental role for affective behavior [2], being involved with goals, intentions and decision strategies. Recent approaches to cognitive systems [3–6] emphasize a strong relation between emotions and agent reasoning, ascribing to affective states a critical aspect in the processes through which agents detect, classify and pro-actively respond to changes in the environment. [1] relates expectations to affective states like surprise, relief, disappointment and hope: by defining surprise as the experienced mismatch between 'what is expected' and 'what is perceived' (at a given level of representation), expectations become a fundamental element for surprise. Directly coupled with

* The research of this paper has been funded by the EU Project MindRACES within the Cognitive Systems area: Contract Number: FP6-511931. (www.mindRaces.org).

expectations, at various levels, affective states can be modeled as control states [6], and used to address reasoning processes, to assess alternatives and choices, and to direct cognitive resources towards anticipatory behavior [4].

In the context of goal-driven activities, expectations processing elicit a twofold contribute: it can be associated with fast processes, either to moderate resources (pro-active means-end reasoning, ability to quickly allocate resources to face environment changes) or to govern slow reasoning as goal selection (making decisions with limited information and bounded resources). The goal of this paper is to integrate the various agent processes required to manage these different roles played by expectations. We refer to Beliefs-Desire-Intention (BDI) and *Goal-Directed* agents, distinguishing between *high level expectation* and *low level expectation*.

High level Expectations are explicit expectations, consisting of predictions about the consequences of decisions, that can be associated with agent's terminal goals and scrutinized at a deliberative level to decide between alternative courses of actions. We refer to the suggestions of Subjective Expected Utilities (SEUs) [7], that have been placed in decision theoretic account as a function of the agents' *beliefs* and *desires* [8, 9]. To build high level expectations we refer to the following independent dimensions: *Belief strength*, as a degree of subjective certainty (the agent is more or less self-confident about their likelihood) and *Goal value*, a subjective importance strictly dependent on their motivating forces, context conditions and mental attitudes (i.e. desirability). This makes it possible to endow explicit expectations with their *valence* according to the agent's subjective purposes. They can be considered *positive* (or *negative*) according to their contribution (or determent) to the ongoing intentions and mental states (e.g. Goals, Beliefs). Our model builds on top of a BDI engine an *expectation-driven* decision making, thus combining deliberative, logical aspects of a BDI model with more quantitative, numerical aspects of decision theory.

Low level Expectations. In dynamic environments, agents should be sufficiently prompt to adapt to evolving contexts, reconsidering their intentions in order to exploit opportunities or avoid damages. Recent approaches to cognitive systems [5] based these capabilities on the appraisal theory [10], ascribing to the effects of emotions the pivotal function of *interrupting practical reasoning when unexpected events require servicing*. Affective coping strategies may be modeled as a momentary interruption of deliberative and practical reasoning processes (e.g. diverting attention to past episodes or focusing sensors to a restricted class of features). Typically these activities are not part of the specification of an agent as for his purposive behavior, rather, can be let to *emerge* as a result of the interactions of various components in the agent's control system. To enhance adaptiveness towards uncertainty, context awareness and to enable agents to respond to urgent needs, we deal with low level expectations by defining *Mental States* (MS) as moods, attitudes and biases the agent adopts to face world changes. The emergent role of MS enables agent to adopt an internal frame where both expectations and emotions are conveyed to inform reasoning, redirecting resources and adapting strategies.

2 Scenario description

To characterize cognitive agents endowed with different kinds of expectations we designed a simulated scenario, where agents are engaged in a foraging task and move in a continuous 2D land map where walls, obstacles and doors delimit rooms, corridors and pathways. The environment holds simulated time and guarantees consistency for entities, artifacts and world objects. Three locations of interest (LOI) present symbolic reference points where three kinds of food can appear, with fixable frequencies. Each class of food has a modifiable *score* and a certain likelihood to appear close to the corresponding LOI. Food scores can have a dynamic value according to a sinusoidal progress (due to simulated *seasons*). Agents have terminal goal of foraging food, composed by a workflow of actions: *1) Look for Food* with the (supposed) best reward; *2) Go to the identified Food item location* and *pick it up. 3) Transport food item*(one at a time) from the original location to the repository and *deposit it.* By releasing food in the repository, agents obtain a reward augmenting their energy. The belief base is built upon a shared ontology of world's objects and artifacts. Navigation capabilities are given with a repertoire of paths (defined as a list of locations to pass through) used to routinize the crossing of rooms and LOIs. Agents are characterized by the following tuple of dynamic, internal resources: $Ag = \langle En, r, Sr, s\rangle$, En indicating the current amount of energy, r the range of vision where sensors can retrieve data, Sr the sensor sample rate, and s the instant speed. We assume that agents burn energy according to a combination of the previous resource costs (e.g. the higher the speed and sensor-rate, the higher the energy cost). Along with the presented workflow, agents can run up against dangerous entities: *fires* can rise with higher frequencies in given dangerous areas and exact agents to experience a short-term reaction (actions and speed are constrained and further costs in terms of energy have to be paid).

3 Agent model

The agent model is built on the top of the *Jadex* engine [3], a multi-threaded BDI framework leading to loosely coupled Beliefs, Goal and Plans, including their mutual relations. Mechanism of deliberation is driven by the evaluation of logic formulae (put in form of Belief formulae) and inhibition links between goals, used to dynamically resolve their mutual priorities. For simplicity, when an entity is sensed within the sensor range, its symbolic description is provided by a preceptor filter and then used for belief update (Fig. 1a).

3.1 Mental States

Once particular events are recognized they may activate background (tacit, passive) expectations and awareness of local contexts. At any instant of time MS

[3] See http://sourceforge.net/projects/jadex/

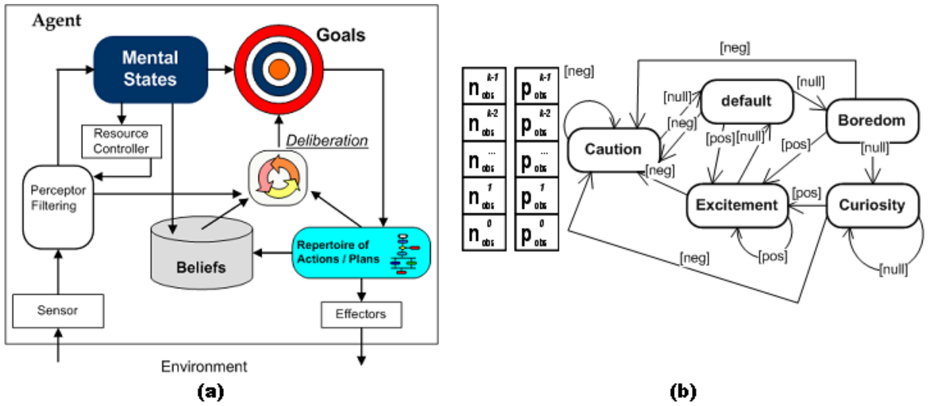

Fig. 1. Agent's architecture (a) and controller for Mental States (b).

agents sense the world around them and store surprising events adding items to the working memory. Here surprise is referred to low level expectations and arises when the agent relieves a mismatch from an unexpected input. For each surprising events agent store an event report in an associative memory: their symbolic representation include time-stamp, positive or negative valence of the originating event, location where the event has been detected and other specialized fields[4]. The content of the associative memory is constantly monitored by an appraisal process that balance the presence of event items and then decide which is the MS to adopt.

Similarly to [11], background expectations are elicited on the basis of the agent's knowledge and are encoded in scripts and frames: we introduce MS as clusters of affective responses to these background expectations placed in form of control states. The current MS is calculated by a transition function. We use two distinct buffers to store positive and negative events. By balancing the presence of event items the appraisal process distinguishes between positive and negative low level expectations. As in [4, 9] the internal model has been described through a push down automaton (Fig. 1b). By using a functional approach, we defined some important roles that affective states play for situated adaptation to contexts and environment dynamism. They trigger a process for a situated intention reconsideration based on rules that: 1) *appraise the events* that interrupt the ongoing actions in terms of positive or negative outcomes and 2) *modulate resources* and control strategies to inform action selection and decision process. For instance, registering a close series of harmful (unexpected) events may elicit the negative expectation that the agent is in a dangerous area and

[4] The valence for observed events is taken in respect to the agent's ongoing purposes and goal, i.e. negative events thwart agents' goals and positive event are those that create new opportunities for goal achievement. In the case of the scenario described above, we distinguished negative events as harmful entities, fire threats, and positive events as food objects and LOI discovering. For more details see [4].

induce to pass to a **Cautious state**(see Fig. 1b). This negative, low level expectation elicit the agent to rise prudence and causes changes both in the long and the short term: it induces alert by modulating attentive resources (i.e. enhancing Sr, augmenting r and reducing speed s), to better check the situation while and before moving (prudence against risks); then it introduces careful in performing dangerous actions by augmenting the control (i.e. performing more perceptive iterations) and/or performing the action in a less risky way, using safest alternatives in repertoires. Positive events induce the agent to **Exciting state** that arouses the agent, increases epistemic activities and the utility of those unexpected positive side effects. The lack of surprises reduces monitoring and increases speed. In the long run, the lack of surprises produces a special mood: **Boredom**. The persistence of boredom can lead to **Curiosity**, whose outcome is to shift from exploitation to exploration attitudes. The agent in this case activates the epistemic goal of exploring and searching for facts in order to improve knowledge and update expectations.

3.2 Affective Expected Utilities

Belief strength can be managed by predictor mechanisms and fully represented within the internal state in the domain of probability as belief in the future with a certain likelihood. By associating a fully represented utility to the different goals, the agent couples each prediction with its desirability and can choose the best among the expected ones. Along with the foraging tasks, a working memory stores information about food quality (reward on goal achievement and food type) and quantity (frequencies, time-stamps and location of the new food added to the belief base). The agent associates a fully represented SEU [7] with each LOI, given multiplying the *Goal value* $Goal_{Rew}$, as an evaluation of the expected outcomes in terms of subjective utility determined averaging rewards stored in a k-length memory buffer for the last k delivered food items (e.g. $U(O_{G_i})$), and a *Belief strength* Bel_{FLOI}, as subjective degree of certainty in domain of probability indicating the likelihood to discover foods near the respective LOI (e.g. $P(O_{a_j}|a_j)$). Once a food is located close to a certain LOI_x, the agent reinforces Bel_{FLOI_x}. Otherwise, when a LOI is visited and no foods are located, Bel_{FLOI_x} is decreased. Agent meta-level-reasoning (i.e. a meta-plan) chooses the alternative by comparing SEUs and selecting the specific plan towards the best expected LOI according to a ε-greedy strategy. Agents are also structurally designed to appraise mismatches: through a feedback signal, results of the purposive action of depositing food are used to reinforce beliefs strength (Bel_{FLOI}) and goal value ($Goal_{Rew}$) and the associated reward expectation.

An integrated *Emotivector* mechanism [12] enables an affective appraisal of the mismatches. The Emotivector is a salience module that is fed with values from a monitored signal, and along with the signal it makes a prediction of its next state and, accordingly, sends extra information along with the signal. This information is based on the appraised discrepancies between the sensed and the predicted value and is composed of two components. An *exogenous* component can appraise a measure of surprise that is proportional to the error made in the prediction. And a further *endogenous* component measures the progress towards a desired value for the signal (*search value*) which generates an affective *sensation*. Agents integrating an Emotivector use a modified model for predictions enabling appraisal of SEUs mismatches.

Predictor: For each food type, the Emotivector monitors the food score, which is its input signal. Prediction is based on a linear model that extrapolates the score relying on the last sensed value considering the velocity of score progression (i.e. first derivative). If the score varies in an approximately locally linear fashion, the linear predictor better captures its dynamics.

Appraisal and sensations: Comparing the expected reward and the achieved reward for a given food score, the different elicited *sensations* give affective measure to the signal. Five sensations can be elicited (Tab. 1). The positive increase

Table 1. Emotivector's sensation description. Each sensation is a subjective and affective appraisal of a goal achievement.

Surprise (S): no expectation on reward, no search value is given.	
Positive increase (S+): a reward is predicted and the effective reward is stronger. Can be related to **excitement**.	*Negative increase (S−)*: more punishment than expected. Can be related to **distress** or **strong disappointment**.
Positive reduction ($+): less reward than expected. Can be related to **disappointment**.	*Negative reduction ($−)*: less punishment than expected. Can be related to **relief**.

($S+$) and the negative reduction ($−) sensations of the monitored signal give a positive indication about the progression of the food score. Hence, when associated with a specific food type, they present a *positive feeling* towards the expectations for that food score. Contrarily, the negative increase ($S−$) and the positive reduction ($+) sensations cause the agent to have a *negative feeling*.

As formally described in [9], emotions may affect the terms of a decision by influencing its desirability and its utility. In our model, agents make use of these valenced feelings to give more or less preference to a certain food type at the deliberative level. This is done by reinforcing the SEU of that food type in the case of a positive feeling, and diminishing it in the case of a negative feeling. We thus provide a further *Affective Bias* (A_b) in order to affect deliberation and high level expectation evaluation. In more detail, the affective agent can use an *Affective Expected Utility* given by:

$$AEU(G_i) = \sum_{a_j \epsilon Plan(Gj)} [A_b \times U(O_{G_i})] \times P(O_{a_j}|a_j) \quad (1)$$

where the utility of a given course of actions is multiplied by an affective bias A_b. This modulating term is positive for positive feelings and negative for negative ones and represents both a qualitative and a quantitative appraisal of the ex-

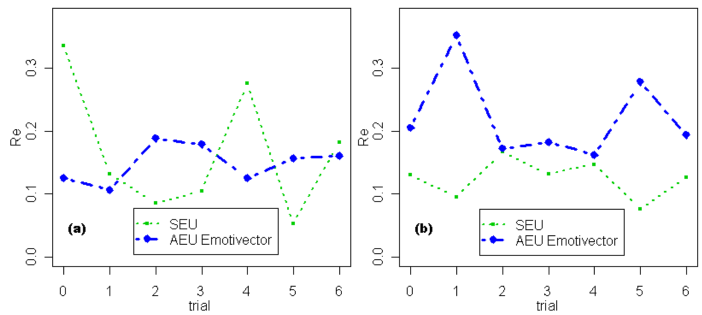

Fig. 2. Comparison between SEU and AEU expectation effectiveness in environment without (a) and with (b) seasons.

perienced mismatch and further introduces a quantitative reinforcement within the deliberation process.

4 Experimental Results

To test how high level expectations are influenced by affective bias, an agent with a decision module based on *SEU* is compared, in the same scenario, with agent using *AEU*. The experiments consisted in multiple trials where the agents were involved with the foraging task. Both agents adopted the same MS module and the same library for navigation, but differed in their deliberation strategy. Each trial had a fixed number of fires (8), a fixed amount of food valuables and a fixed duration (3000 sim-time). The seasonal changes of food scores use a sinusoidal function with different frequencies. As a performance measure, we define *expectation effectiveness* as: $Re = \frac{R_T}{N_T}$, where R_T is the total amount of reward obtained during the single trial and N_T is the total number of food items that appear in the trial. The design, encoding and testing of the integrated system have shown AEU's contribution to performances: for the cases with *static trends of food score* the expectation effectiveness are comparable (Fig. 2a) and agents attain performances of the same order. On the contrary, the AEU agent outperforms the SEU agent enabling seasons and *sinusoidal progress of scores* (Fig. 2b), where, for each tested trial, *Re* for AEU is better. Emotivector's advantages are twofold. First, they assess a better prediction model. The predicted values from the SEU agent are also represented. Relying on a k-length history buffer, SEU prediction is not as adequate for the sinusoidal scores: it is a more conservative prediction, depending more on the past than on the future states, and so it is not suitable for this kind of signal that evolves continuously with time. The second contribution is due to the affective bias: the appraisal of sensations provides the AEU agent with an *modulated motivation* to decide which is the most hopeful area to explore. Notice that SEU agent can't distinguish the case in which a great likelihood is coupled with a low utility, to the case where an high utility is coupled to a scarce likelihood. Differently, AEU agent uses an additional affective element indicating how good is the feeling toward a certain choice, thus anticipating the potential affective consequences of alternatives.

5 Discussion

In this paper we focused on the cognitive role of expectations: they have been assumed to influence goal directed agents at different levels of reasoning, from goal deliberation, to means-end and action selection processes. Expectations have been placed to induce strong relations between agents affective states and their reasoning. We first presented expectations on the low level of the reasoning processes, describing how the agent's appraisal (placed in terms of detection, classification of unexpected events) can elicit mental changes and inform reasoning. Then we provided a decision process for deliberation influenced by affective biases. The new mechanism for affective and deliberative agents improve their anticipatory abilities, and we showed with an experiment its functional advantages. Our results show that in linear dynamism affective bias and linear predictors enable AEU agents to outperform SEU agents. Expectations modeling can be provided in a domain-independent methodology to be exploited in a wide range of applications, from human-machine interactions to multi-agent coordination activities and cognitive models of social interaction like trust, delegation and non verbal communication.

References

1. Castelfranchi, C.: Mind as an Anticipatory Device: for a theory of Expectations. BVAI 2005 (2005) 258–276
2. Ortony, A., Clore, G., Collins, A.: The cognitive structure of emotions. (1988)
3. Ahn, H., Picard, R.W.: Affective cognitive learning and decision making: The role of emotions. In: 18th European Meeting on Cybernetics and Systems Research (EMCSR 2006), Vienna, Austria. (2006)
4. Piunti, M., Castelfranchi, C., Falcone, R.: Surprise as shortcut for Anticipation: clustering Mental States in reasoning. In: Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-07, Hyberadad, India. (2007)
5. Gratch, J., Marsella, S.: A Domain-independent Framework for modelling Emotions. Journal of Cognitive Systems Research **5 (4)** (2004) 269–306
6. Scheutz, M., Sloman, A.: Affect and agent control: Experiments with simple affective states. In: Proceedings of IAT-01. (2001)
7. Savage, L.: The Foundations of Statistics. Dover (1954)
8. Bratman, M., Isreal, D., Pollack, M.: Plans and resource-bounded practical reasoning. Philosophy and AI: Essays at the Interface **4(4)** (1988)
9. Gmytrasiewicz, P., Lisetti, C.: Using decision theory to formalize emotions for multi-agent systems. In: 2nd Workshop on Game Theoretic and Decision Theoretic Agents (ICMAS-2000). (2000)
10. Lazarus, R.: Emotion and Adaptation. Oxford University Press (1991)
11. Schank, R., Abelson, R.: Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures. L. Erlbaum, Hillsdale, NJ (1977)
12. Martinho, C., Paiva, A.: Using anticipation to create believable behaviour. In: Proceedings of the AAAI 2006. (2006)