Expectations driven approach for Situated, Goal-directed Agents

Michele Piunti
Institute of Cognitive Sciences
and Technologies, ISTC - CNR
Università degli studi di Bologna
DEIS - Bologna, Italy
Email: michele.piunti@istc.cnr.it

Cristiano Castelfranchi
Institute of Cognitive Sciences
and Technologies, ISTC
CNR - Roma, Italy
Email: cristiano.castelfranchi@istc.cnr.it

Rino Falcone
Institute of Cognitive Sciences
and Technologies, ISTC
CNR - Roma, Italy
Email: rino.falcone@istc.cnr.it

***Abstract*— Situated agents engaged in open systems continually face with external events requiring adequate services and behavioral responses. In these conditions agents should be able to improve their adaptivity over time, namely 1) to deal with and anticipate relevant changes and critical situations, 2) to temporally define relative priorities between goals varying their importance over time and 3) to use informational feedback to learn from experience and become better at achieving their goals. This work provides an insight to model goal directed agents with these adaptive and anticipatory abilities, based on context awareness and growing experience at achieving their activities. We propose an approach by which affective states are placed as an integrated control mechanism in order to tight different processes and computational modules underlying reasoning.**

I. Introduction

A great variety of goal-directed models of agency, focused at various level on representational, deliberative and action selection mechanisms, have been developed over the last two decades to design adaptive, autonomous and socially interactive agents. We here refer to the goal-directed model of agency, where agents are intended as autonomous, resource bounded entities that attempt to arbitrate between several goals interacting in dynamic, partially observable environments. Typically goal-directed agents are engaged in deliberation to select to which of the concurrent goals devote their resources. Reflecting the original model proposed at the end of 80s [1], traditional deliberative systems process their information reacting in a procedural way: they choose in a repertoire the action to execute according to filtering of conditions (matching rules based on belief formulae, priority hierarchies etc.), whilst the available plan library is handcrafted at design time. An agent can change his environment through a given set of actions and plan operators in order to reach a desired state of affairs from the current state. According to the wide adopted *Belief Desire Intentions* (BDI) model of agency [2], [3], an action is performed when the agent has the intention to achieve a given goal, and the beliefs indicating that the action helps in achieving that goal. Whilst it is theoretically possibly to specify an effective model of behavior in deterministic or probabilistic environments, it is very troublesome to deal in practice with real conditions where the agent has to face with resource limitations (time, computation, memory), partial information (hidden state due to environments constraints, noise, sensor weakness), time-varying goal importance, changing probability distributions and non-stationary, non-probabilistic, non deterministic environments. This negatively reflects upon agent design process and requires the designer to fully understand the dynamics and the complex cases that an agent may face with, and then to implement solutions from scratch for each critical situation may be encountered [4], [5].

The computational model presented here is an attempt to bridge the existing gap between goal directed model of agency and more situated models used in artificial life. Our challenge is to endow deliberative agents with capabilities to operate in *open systems* where dynamism, partial knowledge and unpredictability of future events exact agents to quickly *anticipate* decisions, facing with uncertainty and unexpected events. As a general foundation for adaptiveness in artificial entities [5], there is the need for a proactive adaptation of the internal model over time, in order to exploit informational feedbacks, to learn from experiences and become better at achieving goals. To deal with adaptiveness and anticipation, we refer on a cognitive models of *expectations* as causal precursor of basic *emotions* [6], [7], as far as on recent convergent studies that are pointing out the enhancement of adaptiveness in introducing emotions in reasoning [8], [9], [10], [11].

On the basis of the formal definition of a series of affective states, we provide a description of their functional role, assessing a series of behavioral and mental changes that emotions may induce within agent's internal processes. Our approach is intended at: *1)* Exerting a top-down modulation of emotional reasoning as a result of deliberative process and adaptive responses to relevant events and *2)* Integrating adaptiveness in decision making along with expectations and their causal relation with appraisal/evaluation of events.

The remind of this work is organized as follows: in section 2 we present a model for active and surprise driven perception and belief update, in section 3 we describe an expectation based approach for decisions affected by emotions, in section 4 we present a model for situated reasoning enabling appraisal and coping strategies to unexpected events, section 5 shows as a long term effect of metal states can be integrated in decisions, in section 6 we conclude presenting some related works and providing a final discussion.

II. From expectations and active perception to surprise

Among all the activities an agent may perform during his tasks, we identify two main typologies:

1. Purposive behavior, supported by practical reasoning, is aimed at achieving terminal goals through the use of practical actions and plans. A goal directed agent should use informational feedback to learn from experience and become better at achieving goals.

2. Situated behavior, supporting coping strategies, is aimed at recruiting resources when some *unexpected* event require services. A situated agent should re-define relative priorities between goals varying their importance over time.

Therefore we here identify two *integrated* levels of reasoning, involving cognitive, slow deliberative processes as well as fast automatic and associative ones. Both levels integrate various mechanisms required to manage expectations, used either to assess alternatives and choices, and to direct cognitive resources towards anticipated events. In more detail, we distinguish between *high level, active expectations* and *background, passive expectations*.

At an high level we deal with *explicit expectations* modulating decisions and thus goal deliberation: we include in the reasoning process a quantitative influence on the terms given by the expected utilities used for arbitrating between alternative courses of actions [9]. As we will show later, these influences can be adjusted on the basis of affective appraisal and experiences and allow agent to learn from experiences.

Besides, in order to enhance agent's adaptiveness, we model situated, *background expectations* to elicit a loss of control of certain deliberative processes and to reconsider the course of agent activities. Typically, these particular kind of reasoning is not part of the specification of an agent in his purposive behavior, rather can be let to *emerge* as a result of the interactions in his environments. In [12], Lazarus indicated this particular process as a reflexive re-assessment of the internal state under context awareness, rather than an explicitly deliberated process.

Each of the aforementioned activities should be supported by an adequate perceptive process. Traditional agent architectures use simplified approaches for perception (i.e. based on hard coded rules for controlling sensor apparatus) thus resulting in monolithic and domain specific mechanisms. On the contrary, many evidences pointed out the need for a more abstract, pro-active and goal-driven model of perception [13], [14]. Unfortunetely, active perception is computationally expensive for resource bounded agents. According to the principle of minimal rationality [4], perception filtering should not overload agent processes with continuous belief revision. This exacts a less demanding strategy of reasoning from precepts. To this end, a lighter peripheral filtering can be used only when *relevant* information comes, even while the agent is not actively searching for it. The agent can thus ignore all incoming inputs which are not relevant with respect to the current task and only consider those information which are relevant [15]. In so doing, we identify two different perceptive strategies that an agent may adopt.

1. Purposive tests relate to *active perception* and used for belief update and control of purposive behavior. They are aimed at signalling discrepancies between what is perceived and what is expected in terms of high level, active expectation upon terminal goals achievement. In this case, the agent actively observes the fulfillment of his purposive actions trying to confirm the validity of related expectations.

2. Situated tests relate to *background perception*. In order to gather information of the near contexts, situated agents need the ability to deal with unexpected events, namely events that are not directly and actively under the focus of attention, but that can strongly influence its course.

Many emotions are in tight relation with perception and expectations. For example, surprise is conceived as an expectation based phenomenon: it has been given in terms of a felt signal which provokes an immediate reaction/response of alert and arousal due to an inconsistency (discrepancy, mismatch, non-assimilation, lack of integration) between an incoming input and prior expectations [16], [17]. Surprise has been related to many effects aimed at solving the inconsistency and at preventing possible dangers. Surprise strongly affects attentive processes [18], while [15] show operational advantages of a expectation-driven perception filtering for belief update. We here refer to a particular form of surprise, due to an experienced *mismatch* between a perceived fact and a scrutinized expectation. A specific kind of mismatch-based surprise can be associated to each type of expectation and to each kind of attentive processes. We here identify two different kind of surprise: the former is based on synchronous mismatches appraised upon action completion (i.e. on goal achievement) the latter when passively expected events occurs asynchronously *during* practical behavior (i.e. action execution)[1]. Bringing perception and expectations at the same level of representation, a computational system can detect and quantitatively evaluate the mismatches [18], [17], [15]. In the next sections we deal with expectations and surprise-related behavior influencing goal directed agents at different levels of reasoning. On an higher level, expectations help to take decisions between alternative courses of actions. As for situated cognition, surprise based on passive expectations can enhance context awareness and elicit responses to recruit operative resources to respond *in advance* to changes.

III. Expectations and decision making

Current BDI oriented implementations provide mechanisms for deliberation (goal selection and intention filtering) but don't share common strategies for decision making. Our model builds on top of a BDI engine an *expectation-driven* decision making process, thus combining deliberative, logical aspects of a BDI model with more quantitative, numerical aspects of

[1]Deeper forms of surprise rely on deducibility [18] plausibility [17] of the incoming percept inferred on the basis of the prior knowledge.

decision theory. We identify slower forms of reasoning with high level cognition, decisions between alternative courses of actions used by agents to arbitrate goal selection. To allow agents to take decisions based on the related expectations we model a long term memory entertaining endogenous anticipatory representations. Each (sub)goal is given along with the representation of its activation formulae (typically first-order belief formulae [19]) and a network of inhibition links (indicating if a given goal has the priority on another goal and under which conditions this priority is applicable [20]). Filtering can be managed through a dynamic arbitration network, providing disambiguation between the precondition rules and the relative dependencies (inhibition links) between the concurrent goals. A deliberation engine reacts to changes in the belief base (i.e. internal events thrown by a belief update) and uses the current internal state to filter out enabling conditions for arbitrating the goal adoption.

As for the decision theoretic paradigm of ‘rationality’ [21], [22], an artificial agent may act in order to maximize the expected utility, given multiplying utilities (desirability) and probabilities (likelihood). In our model this strategy is delivered at a meta-level reasoning, typically when the agent has to select between alternatives to achieve mutually exclusive sub-goals.

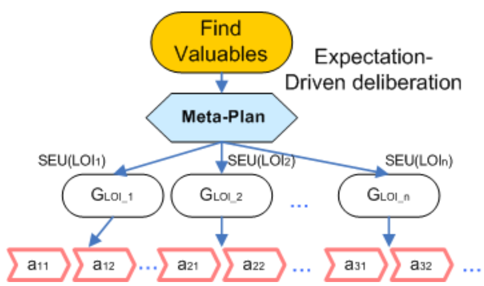

Fig. 1. Given a terminal goal, *Expectation Driven Deliberation* compares Subjective Expected Utilities to choose the most promising course of actions

Imagine an agent being engaged in a foraging task. In normal conditions, the terminal goal is to look for valuables moving to a series of rooms towards some Location of Interest(LOI). Expectation-driven deliberation allows the agent to decide on which LOI to look for considering how the various alternatives are ‘promising’ (Fig. 1). The scrutinized expectations are built upon two independent quantitative dimensions: *Belief strength*, as a degree of subjective certainty placed in terms of likelihood (the agent is more or less certain about their content) and *Goal value*, a subjective importance strictly dependent on desirability of the goal state and the related motivating forces, but also on context conditions and mental attitudes [6], [7]. Given this, *Subjective Expected Utility* (SEU) can be placed as:

$$SEU(G_i) = \sum_{a_j \epsilon Plan(G_i)} U(O_{G_i}) P(O_{a_j}|a_j) \qquad (1)$$

where G_i is the i^{th} goal to adopt between candidates, O_{G_i} is its related outcome, $U(O_{G_i})$ is the subjective utility of that outcome, a_j the j^{th} action of the plan triggered by G_i and $P(O_{a_j}|a_j)$ is the probability of that outcome, given that the j^{th} action of the plan will have the proper O_{a_j} outcome.

Utilities are coupled to rewards obtained upon goal completion and quantitatively assessed in relation to past experiences. $U(O_{G_i})$ is calculated according to the extent to which an intention (i.e. a given sequence of actions) has fulfilled a subjective desire ${O_{G_i}}$[2]. This makes it possible to endow expectations with their *valence*: expectations can be considered *positive* (or *negative*) according to their contribution (or determent) to the ongoing intentions and mental states (e.g. Goals, Beliefs).

Likelihood are subjectively assessed as predictions through a forward model mechanism. In the actual implementation, we are testing different mechanisms for unsupervised learning to determine conditional probabilities of future events, given a sufficiently wide open knowledge base (i.e. EM algorithms for Bayesian networks [23]). Feedback of mismatches between expectations and experienced outcomes are then used to adjust either utilities and predictions. In so doing, even in the same environment, different agents build different subjective models based on their past experiences, thus resulting with different epistemic and motivational states.

A. Emotions modulating high level expectations

Among the consequences of scrutinizing an expectation there is the increment of epistemic activities, aimed at acquiring information from environment *to know whether the expectation can be validated or disconfirmed* [24], [25], [7]. As mentioned, this mechanism is at the basis of any mismatch-based surprise. The idea behind the modulation of expectations with emotions is that an agent can affect the desirability of an outcome by introducing an additional motivation based on an anticipatory feeling[3]: given an active expectation upon a possible reward, agents can appraise their experiences comparing the expected utility and the effective achieved reward. Six cases of mismatch are possible:

1) *Positive increase* ($S+$): the achieved reward is stronger than the one expected. Can be related to **excitement**.
2) *Negative increase* ($S-$): the punishment is greater than expected. Can be related to **distress** or **strong disappointment**.
3) *Positive reduction* ($\$+$): the agent achieve less reward than the one expected. Can be related to **disappointment**.
4) *Negative reduction* ($\$-$): less punishment than expected. Can be related to **relief**.
5) *No Surprise* (NS): goal reward matches the expectation and is exactly the one expected.

[2]In more details, for any given goal G_i, the agent stores achieved rewards and calculates $U(O_{G_i})$ by inferring the next value based on an average, or on a linear projection.

[3]In the field of decision theory, a similar solution has been formally proposed by Gmytrasiewicz and Lisetti [9], while Busemeyer et al. [26] formalized how needs change over time under the pressure of external stimulation and internal deprivation

6) *Surprise due to ignorance (IS)*: the reward is not deducible from prior knowledge due to lack of experiences.

Once appraised, agent can use these feelings to give more or less preference to a certain alternative. The agent may introduce an *affective bias* providing an intrinsic anticipatory effect (the experienced surprise enhances the importance of a certain goal, hence the agent believe to obtain more value from its achievement). We define an *Affective Expected Utilities* (AEU) in terms of:

$$AEU(G_i) = \sum_{a_j \epsilon Plan(G_i)} [A_b \times U(O_{G_i})] \times P(O_{a_j}|a_j) \quad (2)$$

where, respect to the SEU given in (1), A_b represents a qualitative and a quantitative appraisal of the experienced mismatch. It introduces an additional, quantitative reinforcement into the deliberation process and further modulates the expected utility in affective terms. The positive increase ($S+$) and the negative reduction ($\$-$) of the monitored signal give a positive indication about the progression of the goal value. Hence, when associated with a specific decision, they present a *positive feeling* towards the related outcome. Contrarily, the negative increase ($S-$) and the positive reduction ($\$+$) cause the agent to experience a *negative feeling* towards that choice, thus inhibiting its value. This is implemented by reinforcing the utility of a choice with an additional factor, in case of a positive feeling, and diminishing it in the case of a negative feeling. A_b is positive for positive feelings and negative for negative ones:

$$A_b(G_i) = \begin{cases} 0.0 & \text{if } E_s(G_i) \text{ is in } \{NS, IS\} \\ (\gamma_+) * E_r(G_i) & \text{if } E_s(G_i) \text{ is a pos. feeling } \in \{S+, \$-\} \\ (\gamma_-) * E_r(G_i) & \text{if } E_s(G_i) \text{ is a neg. feeling } \in \{S-, \$+\} \end{cases}$$

where $E_s(G_i)$ comes from the last appraised mismatch on G_i's reward, $E_r(G_i)$ is the distance between expected reward and sensed reward, γ_+ and γ_- are discount factors (with $\gamma_+ \ll \gamma_-$).

IV. Background expectations and situated reasoning

A central claim of appraisal theory is that emotions are associated with subjective judgments for the significance of external events (e.g. was the event expected in terms of prior beliefs? is the event congruent with adopted goals? is there the power to alter the consequences of the event?). As shown above, background appraisal allows particular contexts and events to be recognized in order to activate background (tacit, passive) expectations. Agent's *situated perception* envisages causal interpretation of situated events by filtering their features into percepts. The events can then be compared with agent goals and endogenously valued as *positive* (indicating that some event establishes the preconditions for achieving goals or create a new opportunity) or *negative* (some event represent a threat or thwart agent current goals). The idea behind the situated control is that clusters of different coping responses can be arranged around how a situation is appraised. Adopting the model of situated reasoning, coping strategies elicited by different kind of *surprise* can be modeled as a *momentary interruption* of deliberative and practical reasoning processes, e.g. diverting attention to past episodes or focusing sensors and effectors to a restricted area.

A. Mental States and affective control

Stored in an associative memory, noticeable events can be exploited to infer local environment features and activate a passive form of expectations. In so doing, an event that is supposed to thwart an active goal is assessed as a potential undesirable (negative) item, hence the agent reconsider his intentions trying to adopt an alternative action to avoid a threatful state. Otherwise, in case of a positive event, the agent can reconsider his intention in order to exploit an opportunity, or to maintain a desirable state. At any instant of time, agent's situated perception filters the world and store surprising events adding items to a Situated Associative Memory (`SAM`). In this case surprise is referred to passive expectations and arises when the agent relieves a mismatch from an unexpected input coming from the the situated perceptual component. For each of these surprising events, the agent stores in the `SAM` a perceptual report. Reports contain descriptions of a defined set of situated properties: they have a symbolic representation including time-stamp, positive or negative valence of the originating event, location where the event has been detected and other specialized fields[4]:

```
evItem  {   valence: enum value="pos/neg"
            time-stamp: class="Time"
            location: class="Location"
            helps: class="Goal"
            thwarts: class="Goal"
        }
```

Once events are translated to their symbolic representation and stored in the `SAM`, they can be manipulated as percepts. Items have a propositional content but a different nature respect to the beliefs[5]. They are designed to provide both *episodic* and *semantic* contents. The `SAM` is episodic and allows the agent to cache a raw description of situated events, thus enabling the reasoning process to exploit local environment features. The percepts are exploited as a 'fast' source of information to adapt the behavior in the near future and anticipate world changes. The presence of a time stamp for each item ensures to relate each percept to a given time and allow the content of the `SAM` to remain ordered at least for a given field[6]. Besides, the memory is semantic and provides a fast belief base to be handled to infer local environments features. The intuition behind the mechanisms is provided by the well known principle of spatial and temporal locality,

[4] In the case of the foraging scenario, also described in [27], we distinguished negative events as harmful collisions, fire threats, and positive events as food objects, valuables and LOI discovering.

[5] Notice that situated percepts may hold to deceitful appearances [13] including false positive or negative items.

[6] Given the items $I1, I2$ and $I3$, by selecting the time stamps we define an ordered set upon SAM: 1) Either $I1 \preccurlyeq I2$ or $I2 \preccurlyeq I1$ (completeness); 2) If $I1 \preccurlyeq I2$, then $I2 \not\preccurlyeq I1$ (consistency); 3) If $I1 \preccurlyeq I2$ and $I2 \preccurlyeq I3$, then $I1 \preccurlyeq I3$ (transitivity).

according to which one may assess that recently cached items of a certain class are likely to be retrieved in the near future. The amount of item locally present in the SAM can be used as an indication to infer passive expectations about the local context[7].

Transition Function and Information Fusion. SAM's content is constantly monitored by an appraisal process in order balance the presence of items and thus decide which is the MS to adopt. Passing from one state another depends on how the events are relieved and appraised in real time. This process can be described through a push down automaton [27], [9]. Generally the agent supervises the buffers (through a background process) by balancing their registered contents: prevalence of negative items leads to passive expectation of undesirable states (i.e. contingencies, risks), hence to cautious attitudes, while positive events lead to positive expectations (i.e. opportunities) and excitation (Fig. 2). In more details, the current state is inferred by the previous state and the perceived input with a transition function $MsTrans : MS \times IN^{\star} \longrightarrow MS$, where MS is the set of definite mental states and $IN^{\star}$ the input events stored in the (possibly empty) SAM. $MsTrans$ realize an *information fusion* within the symbolic items. Notice that the presence of items of different nature may elicit inconsistencies to be resolved (i.e. presence of elements of different meaning as, for instance, interleaved sequences of positive and negative events). To address this problem $MsTrans$ uses a set of rules for combining and aggregating the items of the same type and circumvent the inconsistencies on the basis of the temporal sequencing given by the time stamps. As suggested by [29], the rules used to govern the fusion can be composed of meta-level and domain specific information. For instance, a simple rule of *balancing* may assert to aggregate the items of a given typology, in order to circumvent the set of lower cardinality and to take into account only the information related to the bigger aggregate. By balancing the presence of items of a given type, the appraisal process suitably distinguishes between positive and negative expectations. A similar approach was used in [27], where two buffers are handled to store positive and negative events and the current sate is let to emerge on the basis of the comparison of buffers sizes (Fig. 2). To prevent the agent to switch to an inconsistent state, the transition function is built to take into account a certain grade of *inertia*, thus providing more robustness against occasional events, false positive or negative items (i.e. due to noise or sensor faults etc.).

B. Functional description

On the basis of a principle of *Analogy*, one may asses that an agent can predict with reasonable accuracy what actions and changes to perform in the near future based on his recent experiences and on the appraisal of local events. In so doing, responses and coping strategies given to a given set of related events can be classified an re-used in analogous situations. Hence, library of coping strategies, action alternatives and resource allocation policies can be clustered within a discrete set of frames used as control states. Effects of coping can

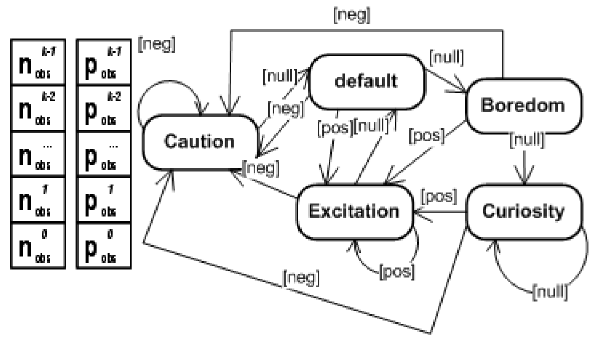

Fig. 2. Controller for Mental States: appraised positive (*p*) and negative (*n*) events are fed to a transition function in order to shift from different mental states.

be modeled in different temporal scale, from immediate and short term reactions, to most persistent long term effects. Given in functional terms, coping strategies includes emotional responses to overturn (in the case of negative emotions) or trigger (in the case of positive ones) control signals to be signalled to the reasoning process. Part of the effect of these signals are conative: mental states, as particular aggregates of control strategies, are modeled to activate particular goals. Besides, MSs are suitable control mechanisms for intention reconsideration. They embed a particular kind of goal activation, bypassing the underlying deliberation processes normally used for practical reasoning. For example, on the short term a MS may attempt to resign the agent to a threat by signalling to the deliberative engine to abandon a goal (thus a related intention) that is becoming inconsistent with the actual belief base or the actual environment state. On the contrary, positive events may elicit goal activation to exploit new opportunities. Furthermore, each MS adopt a *context dependent* configuration of resources (i.e. vision, speed, perception rate, belief update).

Becoming aware of his context, the agent can dynamically adapt his *control frame* in order to reduce performance payoffs and avoid wasting resources for useless activities. Control frames are characterized by the following tuple of dynamic values: $Cf = \langle En, r, Sr, s, G_s \rangle$, *En* indicating the current amount of energy, *r* the range of vision where sensors can retrieve data, *Sr* the situated perception filtering rate, *s* the instant speed and G_s the situated goal to be activated in order to pro-actively respond to the events to cope[8]. Each frame defines the roles that the related MSs play in situated adaptation to contexts and environment dynamism.

Imagine, in the foraging task presented in section III, that the environment presents some threats for agent activities (i.e. the fires, adversary agents etc.). Once the agent has deliberated the best expected location to explore, through the evaluation of

[7]A Similar approach was used by Schank [28], where expectations are generated on the basis of the agent's knowledge encoded in scripts and frames.

[8]We assume that agents spend energy according to a combination of the previous resource costs (e.g. the higher the speed and perception-rate, the higher the spent energy).

MS	Moods	γ_{MS}	Resources		
			r	S_r	S
Default	Exploitation	1.0	.33	.33	.33
Excitement	Reinforcement	1.3	.275	.275	.45
Caution	Prudence	0.5	.45	.45	.10
Boredom	Exploitation	1.0	.33	.33	.33
Curiosity	Exploration	1.0	.45	.10	.45

TABLE I

MENTAL STATES ELICIT THE ADOPTION OF CONTROL FRAMES FOR MOODS, CONFIDENCES AND RESOURCE ALLOCATION POLICIES

the related AEUs, it may happen he registers a close series of harmful (unexpected) events, i.e. fire collisions (Fig. 3.a). This elicit the negative expectation that the agent is approaching to a dangerous area and thus induce him to pass to a **Cautious state** (Fig. 3.b). This negative, background expectation causes the agent to adopt a new control frame, re-allocating his resources to cope harmful circumstances (Tab.I). Cautiousness causes changes both in the long and the short term: firstly it induces arousing by modulating attentive resources (i.e. enhancing *Sr*, looking ahead and augmenting *r* and reducing *s*, see Tab.I). A risk avoidance goal G_s interrupts the ongoing practical action to escape from threats and accordingly the agents arranges activities to better check the situation. On the long term, cautiousness brings to a watchful mood, by reducing the self confidence on beliefs (γ_{MS}), augmenting the control (e.g. enhancing perceptive iterations *Sr*) and/or performing the action in a less risky way(e.g. using safest alternatives in repertoire). Prevalence of positive surprising

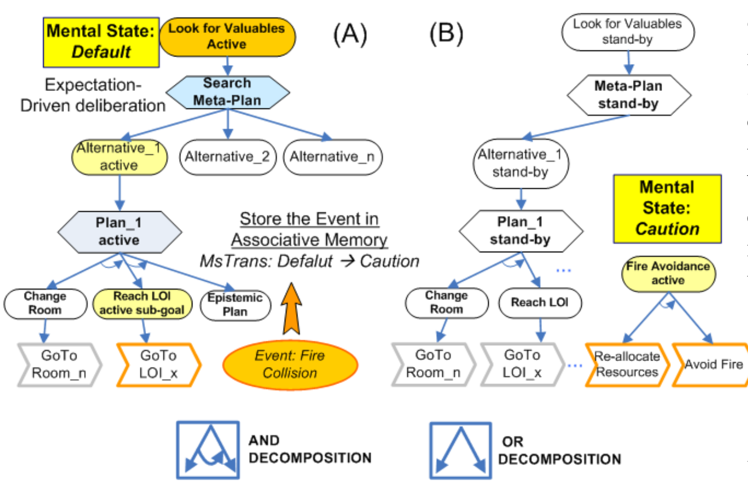

Fig. 3. Intention Reconsideration upon the activation of the Cautiousness Mental State

events induces the agent to shift to **Excitation**, that on the short term is used to arouse the agent, to augment epistemic activities and to search for those 'good' events. A positive surprise (i.e. valuables discovering) may induce the agent to abandon a previous intention and to reformulate his behavior to exploit the new opportunity triggering a new goal G_s. On the long term, excited agent adopt an 'optimistic' mood increasing the confidence (γ_{MS}) of those unexpected, positive events[9]. The lack of surprise progressively empties the SAM and reduces situated perceptive activities. In the long run, it produces a special frame: **Boredom**. Boredom indicates that the environment is almost stationary (no unexpected events are happening) and that the agent can fully exploit his purposive behavior governed by the deliberation driven reasoning. This enhances the subjective confidence in beliefs and in building predictions. Further persistence of boredom leads to **Curiosity**, a control state used to automatically arbitrate from exploitation to exploration activities. The exploration attitude is goal driven: once the agent does not recognize relevant events in his SAM[10], he may infer the low-level expectation that the environment is becoming more static, hence biases his activities towards actions that shows promise to perform a better field coverage and to maintain an updated knowledge. Bypassing the deliberation of practical reasoning, the curious agent pro-actively activates the epistemic G_s of exploring new rooms searching for new facts and events. This has a twofold effect: on the one side it enhances territorial exploration augmenting the chances to discover new LOIs, on the other side it improves knowledge and maintains updated beliefs[11].

V. CONFIDENCE AND MODULATION OF THE PROBABILITY FUNCTION

Effects of MSs can be reconciled with the deliberative level. The intuition behind this integration relies on the fact that each MS endows a certain grade of self-confidence (due to the ongoing mood) that can be related to the belief base. Once we detailed beliefs with a certain strength, one may introduce the self-confidence as a further *discount factor* to affect the likelihood of the predictions. In so doing agents can dynamically adopt a 'more or less confident' capability to build predictions. For example, positive moods provided by excitation can induce the agent to optimistically over-estimate the probability of a certain outcome. On the contrary, negative moods like cautiousness may introduce pessimistic under-estimations. On these basis, respect to the one given in (2), the affective expected utility results:

$$AEU'(G_i) = \sum_{a_j \epsilon Plan(G_i)} [A_b \times U(O_{G_i})] \times [\gamma_{MS} \times P(O_{a_j}|a_j)] \tag{3}$$

where γ_{MS} is associated to the ongoing mental state (Tab.I). By associating a given confidence to the subjective capability to make predictions, γ_{MS} introduces a further affective modulation on agent rationality.

[9]Notice that, differently from Excitement emerging on appraisal of an achieved goal, Excitation has been related to a situated positive surprise.

[10]Heuristic thresholds define the k-length time window used for passing from Boredom to Curiosity.

[11]The benefits of interleaving epistemic and practical activities are generally accepted in situated cognition [24]. Different policies can be retrieved in literature to manage exploitation Vs. exploration. Among others, Ahn and Picard [30] proposed an affective signal to abandon exploitation and trigger the process of exploration.

VI. Discussion

Recent computational models are providing simple affective states in terms of their effects on agent's reasoning, behavior and attentive activities [31]. Their functional roles may enable adaptive and situated behaviors and span from reactive methods of control (similar to those employed in primitive biological organisms [32]) to the control of computational resources [33] and the decision making [22], [11]. In our model we propose a quantitative influence of affective states upon the terms of a rational decision. As in appraisal-inspired models, we provide emotions to coordinate the different computational and physical components required to effectively interact in complex environment [34], [35]. Appraisal based systems like Gratch and Marsella's EMA [35], [11], stresses different relations between emotions and cognition, arguing that emotions are a causal precursor of the mechanisms to detect, classify, and adaptively respond to significant changes of environment. Differently from EMA, we adopted a two step approach. First we distinguished between long-term practical reasoning and situated reasoning. The disambiguation of slow, decisional processes from situated ones elicits a clear methodological separation of concerns and may greatly assist the modeler by breaking down the work into two separate and independent activities: while the former is defined referring to the goal overview and clearly involves decisional processes and deliberation of alternative goals, the latter can be defined through control frames, clustering domain dependent strategies, aggregates of heuristics and functional even affective responses used to respond to local events. In the second phase, we reintegrate the two processes by taking into account the correlations and the relative interactions, enlightening how low situated reasoning can be used to inform higher decisional processes. To this end the contribute of MSs is twofold: from the one side they can relieve the deliberative and the attentive processes from the burdens to process weakly relevant information in decision processes, excluding action alternatives that are likely to be less promising or have vanishing likelihood to be achieved. Besides, MSs provide ready to use action selection and resource allocation policies that may relieve agent's need for resource-demanding and meta decision processes. The emergent nature of affective states enables agent to adopt a mental frame while both expectations and emotions are conveyed to inform reasoning for redirecting resources and adopt long term strategies once a disturbing event is detected.

An additionally effect of modeling mental states is for agent's intention reconsideration. Traditional reconsideration strategies indicate an agent to abandon an intention when a related goal is achieved, when a goal become infeasible or when the agent relieve some inconsistencies between the world state and the external conditions necessary for goal achievement. Our model allows basic emotions to elicit an *interruption* on normal cognitive processes when unexpected events require servicing. Once based on expectations of future states, intention reconsideration becomes anticipatory and can be used to coordinate behavior with prediction of future states.

References

[1] M. Bratman, *Intention, Plans and Pratical Reasoning*. Harvard: Harvard University Press, 1987.
[2] A. Rao and M. Georgeff, "BDI agents: From theory to practice," *Proc. of the 1st conf. on MAS (ICMAS95)*, 1995.
[3] M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge, "The Belief-Desire-Intention Model of Agency," *Proc. of the 5th International Workshop on Intelligent Agents V: Agent Theories, Architectures, and Languages (ATAL-98)*, vol. 1, p. 1555, 1998.
[4] C. Cherniak, *Minimal rationality*. MIT Press, Cambridge, 1986.
[5] P. Maes, "Modeling adaptive autonomous agents," *Artificial Life, I*, vol. (1&2), no. 9, 1994. [Online]. Available: citeseer.ist.psu.edu/maes94modeling.html
[6] M. Miceli and C. Castelfranchi, "The mind and the future. the (negative) power of expectations," *Theory & Psichology*, vol. 12(3), pp. 335–366, 2002.
[7] C. Castelfranchi, "Mind as an Anticipatory Device: for a theory of Expectations," *1st intern. symposium Brain, Vision, and Artificial Intelligence (BVAI 2005)*, pp. 258–276, 2005.
[8] R. W. Picard, *Affective Computing*. MIT Press, 1997.
[9] P. Gmytrasiewicz and C. Lisetti, "Using decision theory to formalize emotions for multi-agent systems," in *2nd Workshop on Game Theoretic and Decision Theoretic Agents (ICMAS-2000)*, 2000.
[10] M. Scheutz, "How to determine the utility of emotions," in *Proc. of AAAI Spring Symposium 2004.*, 2004.
[11] J. Gratch and S. Marsella, "A Domain-independent Framework for modelling Emotions," *Journal of Cognitive Systems Research*, vol. 5 (4), pp. 269–306, 2004.
[12] R. Lazarus, *Emotion and Adaptation*. Oxford University Press, 1991.
[13] J. L. Pollock, "Taking perception seriously," in *Proceedings of the 1st International Conference on Autonomous Agents (Agents'97)*, W. L. Johnson and B. Hayes-Roth, Eds. New York: ACM Press, 1997, pp. 526–527.
[14] D. Weyns, E. Steegmans, and T. Holvoet, "A model for Active Perception in situated Multi-Agent Systems," in *Proc. of the 1st European Workshop on MAS*, 2003.
[15] E. Lorini and M. Piunti, "The benefits of surprise in dynamic environments: from theory to practice," in *Proceedings of 2nd conference on affective computing and intelligent interactions (ACII07)*, 2007.
[16] R. Reisenzein, *The message within: The role of subjective experience in social cognition and behavior*. Psychology Press, 2000, ch. The subjective experience of surprise.
[17] E. Lorini and C. Castelfranchi, "The cognitive structure of surprise: looking for basic principles," *Topoi: an International Review of Philosophy*, vol. 26(1), 2007.
[18] A. Ortony and D. Partridge, "Surprisingness and Expectation failure: What is the difference?" in *Proc. of the 10th Inter. Joint Conf. on Artificial Intelligence*, Los Altos, CA, 1987, pp. 106–108.
[19] J. Thangarajah, L. Padgham, and J. Harland, "Representation and reasoning for goals in BDI agents," in *In Proc. of the 25th Australian Computer Science Conf. (ACS2002)*, 2002.
[20] L. Braubach, A. Pokahr, D. Moldt, and W. Lamersdorf, "Goal representation for bdi agent systems," *The Second International Workshop on Programming Multiagent Systems (PROMAS-2004)*, 2004.
[21] L. Savage, *The Foundations of Statistics*. Dover, 1954.
[22] J. Doyle, "Rationality and its roles in reasoning," *Computational Intelligence*, vol. 8 (2), pp. 376–409, 1992.
[23] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM Algorithm." *Journal of the Royal Statistical Society*, vol. 39(1), p. 138, 1977.
[24] D. Kirsha and P. Maglio, "On distinguishing epistemic from pragmatic action," *Cognitive Science: A Multidisciplinary Journal*, vol. 18(4), pp. 513–549, 1994.
[25] E. Lorini and C. Castelfranchi, "The role of epistemic actions in expectations," in *In Proc. of the ABiALS workshop*, 2004.
[26] J. J.R. Busemeyer, J. J.T. Townsend, and J. Stout, "Motivational underpinnings of utility in decision making: Decision field theory analysis of deprivation and satiation," *Emotional Cognition*, 2002.
[27] M. Piunti, C. Castelfranchi, and R. Falcone, "Surprise as shortcut for Anticipation: clustering Mental States in reasoning," in *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-07)*, Hyberadad, India., 2007.

[28] R. Schank and R. Abelson, *Scripts, Plans, Goals and Understanding: an Inquiry into Human Knowledge Structures*. Hillsdale, NJ: L. Erlbaum, 1977.

[29] L. Cholvy and A. Hunter, "Information fusion in logic: A brief overview," in *ECSQARU-FAPR*, 1997, pp. 86–95. [Online]. Available: citeseer.ist.psu.edu/cholvy97information.html

[30] H. Ahn and R. W. Picard, "Affective cognitive learning and decision making: The role of emotions," in *18th European Meeting on Cybernetics and Systems Research (EMCSR 2006)*, Vienna, Austria., 2006.

[31] M. Scheutz, "Agents with or without emotions?" in *Proceedings of FLAIRS 02, AAAI Press, 89-94.*, 2002.

[32] M. Scheutz and A. Sloman, "Affect and agent control: Experiments with simple affective states," in *Proceedings of IAT-01*, 2001.

[33] H. A. Simon, "Motivational and emotional controls of cognition," *Psychological Review*, vol. 74, pp. 29–39, 1967.

[34] L. Cosmides and J. Tooby, *Handbook of Emotion (Second Edition)*. Guilford Press, 2000, ch. Evolutionary Psychology and the Emotions, pp. 91–115.

[35] J. Gratch and S. Marsella, *Integrated Models of Cognitive Systems*. Oxford University Press, 2006, ch. The architectural role of emotions in cognitive systems.