Surprise as shortcut for Anticipation: clustering Mental States in Reasoning *

Michele Piunti, Cristiano Castelfranchi and Rino Falcone
Institute of Cognitive Sciences and Technologies - CNR
{michele.piunti, c.castelfranchi, r.falcone}@istc.cnr.it

Abstract

To enhance effectiveness in real world applications, autonomous agents have to develop cognitive competencies and anticipatory capabilities. Here we point out their strong liaison with the functional roles of affective mental states as those of human-like metaphor: not only the root elements for both surprise and anticipation are expectations, but also part of the effects of the former elicit efforts on the latter. By analyzing different kinds of expectations, we provide a general architecture enhancing practical reasoning with mental states, describing and empirically evaluating how mental and behavioral attitudes, emerging from mental states, can be applied for augmenting agent reactivity, opportunism and efficacy in terms of anticipation.

1 Introduction

While cognitive systems are attending in growing interest for anticipatory behaviors, multidisciplinary studies remark liaisons between anticipatory mechanisms and functional role of emotions. Otherwise is generally accepted that to enhance human-like effectiveness in real world applications, autonomous agents have to develop higher level cognitive competencies. In this paper we claim and explain how autonomous agents can be anticipatory, able to deal with future events, and how this is important not only for robotic but also for software agents. In particular, we define cognitive anticipatory agents not simply endowed with some statistical learning or prediction, but also with true *expectations*, related with their epistemic states (*Beliefs*) and their motivational states (*Goals*). From a computational viewpoint, we deal with world predictive representations and we refer to expectations that can be framed among internal state and knowledge-base. Although adopting a BDI-like [Rao and Georgeff, 1995] approach, we do not introduce for expectations a new primitive, but we build them on the basis of beliefs and goals. Expectations processing in real time requires monitoring, appraisal, revision and updating, while, along practical reasoning [Bratman *et al.*, 1988], expectations are directly involved at various level in goal deliberation, planning, intention reconsideration, learning and action control.

In addition, expectations have a foundational role in emotion life-cycle and expectation enabled agents are leaning to be 'surprised' according to a human-like behavioral metaphor. Defining surprise as a function of the experienced mismatch between what is expected and the perceived data (at a given level of representation), expectations become "prerequisites" for surprise, thus different kinds of expectations holds to different kinds of surprise. Here we point out that sources of surprise (generally speaking, the "unexpected" signals) can have either negative or positive consequences on purposive behaviors when they are considered in terms of penalties, costs rather than benefits, advantages.

In sections 2 and 3 we give a reformulation of the problem from a cognitive perspective, in terms of *Mental States*. We propose that different properties and outcomes of surprise can be modeled in terms of mental states/attitudes clustering suitable reactions and functional efforts. In particular, we analyze surprise outcomes in autonomous agents engaged in a foraging task in a risky world, where surprise attitudes are significant either to become cautious, prudent, careful in harmful circumstances, either for reinforcing expectations, for enhancing knowledge, for learning and appraisal processes. Some of the functions that surprise plays for an adaptive behavior and an adaptive cognition can be seen in anticipatory terms of opportunistic adjustment to circumstances, of immediate 'reactions', but also in terms of intention reconsideration, attention, belief revision, learning. Specialization of above effects induces balancing of agent resources, introduces pay-offs in performances and may hold to domain dependent decision strategies. Since a theoretical model, we design *Cautiousness*, *Excitement*, *Boredom* and *Curiosity*, giving them the special 'moods' that agent uses to adapt to unexpected chances and to anticipate the world. In section 6 an experiment discussion is proposed to evaluate how effectiveness (in terms of anticipation) is affected both by mental states and environment dynamics. Final discussion is given in section 7.

2 Expectations and Practical Reasoning

Goal directed architectures focalize in deliberation process among set of goals, but let intention making and execution of plans in a functional, even purely reactive form: agents pro-

*Work supported by the EU project **MindRACES**, FP6-511931, www.mindRaces.org.

cess information reacting in a procedural way and choosing in repertoire the plan to execute according to filtering of conditions and belief formulae. No native support was defined for dealing with the future (e.g. future directed intentions), neither prevision models. Even if the idea was to align agent performances to real worlds environments, soon the inadequacy of the model has shown drawbacks. Real world applications face with dynamism, low accessibility of environment state and constraints.

The main problem to be addressed in real-time, situated and goal-directed model of agency is *narrowness-boundedness* in various aspects like computational power, time available to take decisions, knowledge and memory. Originally proposed in practical reasoning by Bratman, intention reconsideration is a costly and binding process: once goal is adopted, agent commits its intention to a given state of affairs, and devote resources for achieving it. Traditionally optimization of intention reconsideration processes relies on the two levels of goal deliberation and plan selection: abstractly, agents should break plan and shift their intention if: *(i)* the related beliefs (context conditions) become false; *(ii)* the committed plan is no longer achievable and there are no alternatives; *(iii)* the root-goal (or the meta-level intention) is inhibited by other goals (or by other meta-intentions). [Kinny and Georgeff, 1991] analyzed different reconsideration strategies according to a 'degree of boldness'[1]. Their experiments showed that *cautious* agents outperforms *bold* agents in cases of high world dynamism but, if the environment is static, all the costs for frequent intention reconsideration are wasted. In a successive work [Kinny *et al.*, 1992] introduced the 'cost of sensing' and showed that agent effectiveness decrease with the increasing of the sensor effort. On the contrary, the optimal sensing rate increase along with the world dynamism. The experiments considered on the one side the time spared by early detection of changes, on the other side the costs for too frequent sensing.

Agents reasoning has to rely on uncomplete knowledge: their beliefs couple with the real world state at various grades of adherence, resulting inconsistent due to ignorance and uncertainty . Unlike completely observable environments, in partially observable environment (POEs) the observation at time t may not provide sufficient information to identify s_t[2].

Further, disambiguation between internal 'drives' or motivations and goals is still an unsolved issue. For instance Boltzmann selection techniques are used for controlling the trade-off between exploration and exploitation; other models consider environmentally mediated model of emotions[3] or some functional role of affective states.[4] Many of these approaches show drawbacks in their cognitive model and generally lack in opportunism against unexpected events: anticipatory agents should have strong proactive capabilities in using uncertain beliefs within deliberation processes, strategies for adaptive intention reconsideration [Kinny and Georgeff, 1991] and sensing, disambiguation between motivations and goal hierarchy.

In the next sections we propose a new approach, based on the high-level role of expectations, eliciting emotional states and attitudes. As we see, this can be reflected in design of reasoning and decision making processes affected by emotional signals.

3 Expectations, Surprise and Anticipation

We refer to anticipation outcomes as agent changes on mental attitudes in decision making due to what is expected. While the association between internal state, processed input and deliberation process is generally defined in design time, anticipatory strategies for intentional (goal driven) behaviors requires agent to build some predictive representations. In [Ortony and Partridge, 1987] cognitive expectations are given in terms of practically deducible propositions, defined as symbolic beliefs that can be fully represented in memory, or logically inferred. Our claim is that expectations in deliberative agents can be coupled with Beliefs and Goals: since these basic components, we attain expectations as a *molecule* expressing *emerging* attitudes , in part epistemic (to check whether the prediction really fits the world, to monitor the success of the behavior) in part motivational (the agent who build expectations is 'concerned' and some of its goal are involved) [Castelfranchi, 2005]. At a cognitive level, we distinguish here between high and low level expectations: $Expectation\alpha$: more explicit, consists of fully represented predictions about decision outcomes and can be associated with alternative courses of actions; $Expectation\beta$: dealing with those expectations with weak level of representation, due to lack of beliefs, uncertainty, ignorance. At a quantitative level, we refer to two independent dimensions:

1. *Belief strength*, as degree of subjective certainty. The agent is more or less sure and committed about their content.
2. *Goal value*, a subjective importance strictly dependent on context conditions and mental attitudes.

[1] They define *bold* agents which never stop to reconsider intentions until the current plan is fully executed, and *cautious* agents which stop to execute and reconsider their intentions after the execution of every action.

[2] Approaches to POE use beliefs state methods (e.g. POMDP) and approximation techniques. In continuous space, intractabilities arise in updating belief and in calculating value function and suddenly agents may operate with hidden states.

[3] [Parunak *et al.*, 2006] endowed a BDI architecture with agent motivational state composed by a vector of seven drives, where the current emotions modulates behavior.

[4] In their computational model of surprise [Macedo and Cardoso, 2004] proposed a solution for exploration of unknown environments with motivational agents where surprise holds to intentions and "action-goal", thus eliciting action-selection through evaluation of utility functions. We guess in this model two different levels are 'collapsed', resulting surprise concurrently a "motivation" (like curiosity and hunger) and a "intention" ("the event with the maximum estimated surprise is selected to be investigated"). As they define agent first goal to enhance knowledge, surprise becomes both a mechanisms for "appraisal of unexpectedness" and a further terminal goal, pursued maximizing expected novelty. In addition, the definition and functional role of curiosity remains unclear respect to the one given for surprise: they identify both as the novelty-discrepancy between the perceived data and the previous beliefs stored in memory.

This kind of expectations already allow agent to experience surprise, elicited by the expectation resulting to be wrong after the fact (expectation invalidation). Thus, expectations become 'prerequisites' for surprise: more precisely surprise is due to (and a signal of) a *mismatch*, an inconsistency between agent expectations (given by previous knowledge) and the incoming information (actual input), compared at a given level of representation [Castelfranchi and Lorini, 2003].

Current computational models lack quantitative description of the functions that surprise play for an anticipatory behavior and an adaptive cognition. Psychoevolutionary model of [Meyer *et al.*, 1997] proposed that surprise-eliciting events holds to *(i)* appraisal of a cognised event as exceeding some threshold of unexpectedness; *(ii)* interruption of ongoing information processing and reallocation of resources to the investigation of the unexpected event; *(iii)* analysis/evaluation of that events; *(iv)* immediate reactions to that event and/or updating or revision of the "old" schemas or beliefs. Recent works in neuroscience and neuroeconomics identify the important role of emotions in decision making and anticipatory goal deliberation processes. It has been showed that affective states help to "decide advantageously before knowing the advantageous strategy" [Bechara *et al.*, 1997]. To improve the speed of learning, regulate the trade-off between exploration and exploitation and learn more efficiently, [Ahn and Picard, 2006] introduced decision making strategies using affective components given by affective anticipatory rewards. They proposed a cognitive system with affective biases of anticipatory reward, modeled by positive and negative valence of evaluation of expectation[5].

We propose that effects of the expectation processing can be seen beginning from the strength of its component (Beliefs and Goals) and can rely on the positive or negative consequences upon agent purposes. The prediction of possible world outcomes triggers a given anticipatory behavior (e.g. aimed at preventing threats). Relatively to goals, we consider the 'goal failure' in terms of excitement or frustration, when the goal is achieved at more or less level than the one expected. Relatively to beliefs, we consider the epistemic capabilities aimed at solving the inconsistency. We also guess that different kinds of expectations hold to different kinds of surprise. Expectation α failure elicit updating of those explicit expectations coupled with goals, reinforcement in expected utility and exploitation of purposive behavior. On the other side, unexpected β signals elicit adaptation, reactiveness, backtracking, intention reconsideration, shift of motivations, opportunism, investment in resources. We distinguish between:

1. *Long-term and short-term effects*, respectively related to decision making, learning and adaptive behavior, arousing, invoking and mobilizing resources attention.
2. *Mental and behavioral changes*, affecting motivations, intention reconsideration, self-confidence, action execution and selection.

[5]In particular, the extrinsic rewards (rising from external goal, or cost) are integrated with intrinsic rewards (rising from emotional circuits), internal drives and motivations: these affective anticipatory rewards are modeled to improve learning and decision making.

Specialization of the above effects induces balancing of agent resources, introduces pay-offs in performances and may hold to domain dependent decision strategies.

4 Foraging in Risky World

To test agent performances so that different strategies can be significantly compared, we design a test-bed scenario. Environment captures features of real world domains, allowing flexible control of the world dynamics. Navigation capabilities are given with a repertoire of paths (defined as list of location to pass through) used to routinize crossing rooms and LOI. Agents move in a continuous 2D land map where walls, obstacles and doors (that can be open or close) delimit rooms, corridors and pathways. Environment holds simulated time and guarantees consistency for entities, artefacts and world objects. Belief base is built upon a shared ontology of world objects and artefacts: three Locations of Interest (LOI) present symbolic reference points where three kind of food appear, with fixable frequency. Each class of food has modifiable 'score' and a likelihood to appear near the corresponding LOI. Agent works for the terminal goal of foraging, composed of the following workflow of actions: *(1)* Look for Food with (supposed) best reward; *(2)* Go to the identified Food location and pick up it; *(3)* Transport Food (one at a time) from the original location to the repository and deposit it. Releasing Foods in the repository, agents obtain a 'reward' augmenting energy, calculated from the basis of the original food score decreased with a *decay factor* straight depending on the transport interval. Decay is introduced to enhance cost in duration of actions. Agents are characterized by the following tuple of dynamic resources:

$$Ag = \langle En, r, Sr, s \rangle \quad (1)$$

En indicating the instant 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 agent burning energy according to the combination of previous resource costs (e.g. the more speed and sensor-rate is high, the more agent will spend energy).

Along the presented workflow, agents can run up against dangerous entities. Fires behave according to a two state lifecycle periodic function: in their first shape they are in a *smoke* 'premonitory' state, then they become real harmful *flame*. At the beginning of each period, fires change their location with discrete movements. Moreover Fires can rise with higher frequency in given dangerous areas. Against fires collisions, agents have to reconsider their intention (e.g. adopting fire avoidance behavior), and they are constrained to experience a short-term reaction: actions and speed are constrained and further costs in term of energy have to be paid. Agents expire when their energy goes to zero[6].

[6]Notice that in order to appreciate cumulative effects of different strategies and behaviors on the long term, no fatal event are introduced.

5 Design

As for the agent's kernel we adopt Jadex engine[7], a multi threaded framework for BDI agents where the architecture leads to a loosely coupled Beliefs, Goal, and Plans, including their mutual relations, through agent descriptor [Braubach *et al.*, 2005]. Jadex deliberation is driven by the evaluation of logic formulae (put in form of Belief formulae) and 'arcs of inhibition' between goals for dynamically resolving their priority. As explained below, to model expectations and effects of surprise at a system level, we pointed out that they take place at various level of reasoning and we modified the original BDI schema introducing α expectation processing (with Expectation Driven Deliberation) and β expectation processing (with Mental States designed for clustering attitudes for adaptive and anticipatory changes).

5.1 Subjective Expected Utility

We relate α expectations with agent epistemic states (Beliefs) and motivational states (Goals). Subjective Expected Utilities (defined in decision-theoretic accounts as a function of the agent's *beliefs* and *desires* [Bratman *et al.*, 1988]) are included at meta level reasoning as mechanisms for Goal deliberation. Along 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 any food added in the belief base). These are the re-evocable traces of previously achieved actions: agents associate to each LOI a SEU value given by Bel_{Rew} (determined averaging rewards stored in a k-length memory for the last k delivered Foods) multiplied with P_{FLOI} (indicating the likelihood to discover foods near the LOI). Through a feedback, actual results of purposive actions of depositing food are compared against expectations: once a food is located, agent will reinforce P_{FLOI} (correct expectation), otherwise, when a LOI is visited and no foods are located, P_{FLOI} will be weakened (wrong prediction). At the meta-level-reasoning, agent chooses to look for $food_x$ at the corresponding LOI_x by comparing their SEUs and by adopting the epistemic goal toward best expected LOI, according to a ϵ-greedy strategy. SEU processing runs in "discrete time", upon goal achievement or at action completion. Identifying with sensors a set of foods, agent adds them to the beliefbase and heads for the nearest one, observing topology, constraints and obstacle bounds. Deliberation from searching to pickup action is triggered when a food is located, from pickup to homing when a valuable is carried. Note that if a nearest food f_j is located, a new intention inhibits the current plan. Further transition to searching is caused when agent achieves the *drop action*.

In addition, means-end reasoning processes are introduced when agents choose a path between the available ones (*means*), to reach a target location (*end*). The *risk* is a negative α expectation (a threat). Quantitatively risk is a fully represented variable inside the paths that agents use to move and is augmented by unexpected negative events (e.g. fire or smoke threat).

[7]See Jadex project at http://sourceforge.net/projects/jadex/

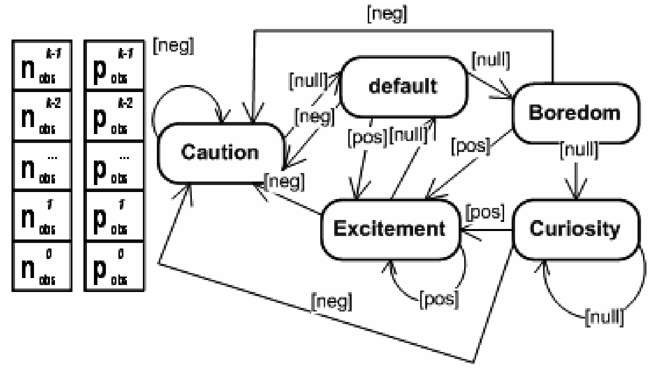

Figure 1: Transition Function for Mental States uses two stacks for storing Positive and Negative events.

5.2 From reactiveness to anticipation

A reactive agent is one that makes decisions about what action to take relying on its current state and on input from sensors. In this case, states can be modeled by a (stochastic) finite state automaton (FSA) where each state represents an action -or a plan- the agent is executing, while transitions get it to other states. Otherwise, generalized Markov process evaluates the past k states: if an agent has instantaneous k-length knowledge of the world, it could dynamically change its internal state, allowing reasoning as a whole to adapt to changes[8]. In most situations, such global knowledge is impractical or costly to collect but for sufficiently slow dynamics, agents can correctly estimate the environment through local observations. This kind of agents can be described by a push down automaton (PDA) where the stack stores the expectation invalidations items. By evaluating the cached items agent can foresee the current environment state (infer $\beta expectations$) and consequently adapt its attitudes. Items have both an *informative* content (e.g. timestamp, Location, event type), and a *semantic* content, because they are coupled with the positive (benefits) or negative (disadvantages) effects that the event entails. Hence, by using two distinct buffers to store these positive and negative events agents, the agent can observe the world in terms of positive opportunities and negative circumstances. Starting from these series of local observations, a background process periodically define the mental state to adopt, through a transition function (see Fig.1).

5.3 Mental States

After a series of positive surprises, unexpected opportunities and helpful novelties, agent may tend to reinforce beliefs: for instance, dropping food with unexpected good reward holds agent to become *Excited*, reinforcing SEU in looking food of that type near the respective LOI. Exciting surprises are internal signal for arousing the agent, for increasing the explorative activity and for searching for those good events. Hav-

[8]A reactive agent can be considered as ordinary Markov process, where future state depends only on its present state. If the memory on the past states has length k, then the agent that is making decisions about future actions rely on the past k states and the process can be represented as a generalized Markov process of order k.

ing registered a close series of harmful events may signify agent is in a dangerous area. The Mental State coupling with alerting (negative) surprise is *Cautiousness*, which gives the anticipatory effort of risk avoidance (i.e. anticipating threats). Being cautious in a risky world, with hidden state, means to become prudent and to adopt safe behaviors. We design agent cautiousness distinguishing two aspects: firstly caution elicits arousal and alert, holds to become more vigilant, to look ahead, to check better while and before moving (prudence against risks); secondly to be careful in doing dangerous actions, either augmenting the control or doing the action in less risky way (e.g using alternatives in repertoire). We suppose cautious agent able to escape from threats using its safest plans. The investment in resources is exploitable for adapting, learning, noticing world regularities and anticipating threats, but introduces pay-offs: [Castelfranchi *et al.*, 2006] showed that mobilizing more resources for epistemic actions (actions explicitly directed to enhance knowledge) and attentive processes has direct effects in reducing promptness and speediness and side effects in bodily reactions, as further energy consumption. Caution holds to behavioral and mental changes: we identify the following activities: *(i)* Perceptive investment, reallocation of attentive resources *S, Sr, v*, looking ahead, updating, focusing; *(ii)* Belief Revision, e.g. signalling a dangerous area, increasing the path risk; *(iii)* Adapting of self-confidence in beliefs (and expectations); *(iv)* Intention reconsideration in the sense of selection of safe actions (e.g. selection of safest path). We design *Boredom* as the mental state coupling with lack of surprising events (i.e. empty buffers). The transition function gets to Boredom by persistence of empty buffers (where persistence is defined by heuristic thresholds). In the long run, further lack of surprise items produces a special 'mood', *Curiosity*, whose outcome is to shift from exploitation to exploration attitudes. In this case the agent activates the new goal of exploring and searching for unexpected event in less visited areas, in order to update knowledge and expectation models. Curious agent uses an alternative searching policy: while the first strategy is based on evaluation of SEU (according to which agent selects searching actions towards best expected LOI) the second has the purpose to 'attain knowledge' in order to update internal world representations, predicting models and beliefs: anticipation is indirectly elicited by appraisal, beliefs and expectation revision. Notice that Curiosity implies abandoning risk

Mental States	**Attitudes**	Resources		
		r	S_r	S
Default	Exploitation	0.73	0.73	0.73
Excitement	Reinforcement	0.6	0.6	1.0
Caution	Prudence	1.0	1.0	0.2
Curiosity	Exploration	1.0	0.2	1.0

Table 1: Mental States and weights for epistemic resources defined in *(1)*. Note that the global amount of resources is limited to 2.2 for each state.

evaluation: it elicits the adoption of a new explicit epistemic goal, leading the agent towards those areas where it foresee to enhance knowledge.

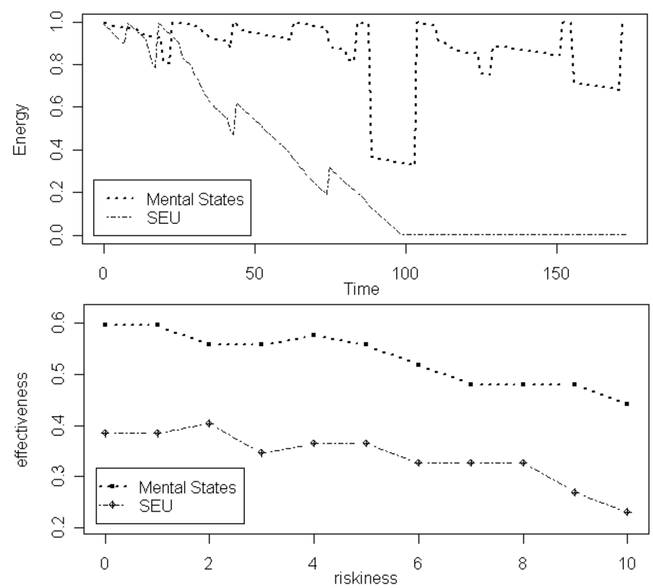

Figure 2: Agent energy (a) and effectiveness (b) comparision

6 Experimental Results

We compared a benchmark agent with subjective expected utility (SEU) and an agent adding to subjective expected utility processing also mental states (MS) module. Experiments have been conducted in environment with different level of riskiness, indicated by the presence of more or less harmful objects (e.g. fires). Our first series of experiments extracted the course of energy in function of time: as in [Castelfranchi *et al.*, 2006], while with low presence of riskiness MS and SEU agents have comparable energy trends, we noticed agent mean energy decay according to enhancing of risks. Fig.2a shows the trial with the presence of 7 fires entities. It is interesting to note the discontinuities in slopes of the functions, caused by transitions of mental states and, consequently, by the different energy consumption due to resource allocation. Peaks indicate goal achievement (food released in the repository) and consequent energy recharge; downfalls indicate damaging by fires. As we expected, MS agent outperformed SEU agent, being able to adapt behavior and reduce harms. SEU agent collided with more fires and this caused it to expire after 1200 sim time. Due to cautiousness, repertoire of safe actions and risk evaluation in choosing of paths, MS agent acted in safeness, succeeded to escape from threats and suffered less damages. In these cases MS agent is able to anticipate dangerous areas and collisions and negative effects are minimized until agent shift to the cautiousness mental state. It also been noticed that agent dynamics introduce payoffs. During their lifetime, SEU achieved 5 forages while MS agent 4: SEU boldness and commitment ensure higher speed and less time to accomplish the task. Otherwise, trials with lower level of riskiness showed better MS agent performances in searching on the long term, due to curiosity. MS agent showed better exploration competencies, when food valuables near LOI are consumed and the usual SEU strategies lacks.

Although the course of energy in function of time gives instantaneous snapshots on the single trial, this metric gives weak contribute to define absolute performances because of dependencies on experiment length, independent random distribution for food, fires and dangerous areas, world dynamism and agent competition for valuables. As in [Kinny and Georgeff, 1991] we defined agent *effectiveness* as the number of achieved task (delivered foods) divided by the total amount of achievable task (total amount of food appearing in the trial). Because of random variations in the environment, effectiveness has a fluctuating course before converging, hence individual trial has to be sufficiently long for the effectiveness to become stable. Measuring effectiveness in function of time, for fixed amount of food, riskiness and world dynamism, we defined the standard trial length of 2100 sim time. We defined agent *characterization* averaging performances of 6 trials. The decreasing curves presented in Fig.2b resulted running characterizations in 10 conditions of growing risk. For both agents, effectiveness comparison shows maximums in the trials with no fires (safe environment). MS agent outperformed SEU agents in all conditions for two main reasons: on the one side cautiousness enable MS agent to avoid fires in risky environment, on the other curiosity allows MS agent to move from static areas, poor of stimulus and valuables, to more dynamics areas, where the likelihood to discover foods -via surprise- is higher. Global effectiveness is 1.0 for low riskiness: in these cases, the agents generally succeeded to collect the whole set of foods. Mean riskiness represents the greatest difference in effectiveness. For greater presence of fires, some foods remain uncollected at the end of the trial either because SEU agents expire before the end of the trial due to recurrent fire collisions, either MS agents pay growing costs for risk avoidance behavior and intention reconsideration.

7 Discussion

This work discussed anticipation through expectations processing and affective efforts in goal driven agents. We examined a cognitive model for agents with mental states affecting reasoning processes; we experienced how related attitudes influence agent performances both on the basis of the environment features and of the cognitive model implemented. By providing a general architecture for intelligent agent clustering attitudes in mental states, we defined flexible domain-dependent decision strategies. Particularly, we showed experiments in environments with growing level of riskiness, where agents with mental states outperforms agents with more traditional utility strategies. We consider these cognitive functions as a fundamental building block for the anticipatory behavior that is the real challenge for the future cognitive software agents and autonomous robots.

This model can be extended in important directions, for instance social activities requiring high level modeling, like trust, delegation and reliance in cooperative or competitive tasks. As for mechanisms, we used PDA for modeling mental states controller, anyway the system is designed for testing alternatives (e.g. salience models), may be resulting more effective.

References

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

[Bechara *et al.*, 1997] A. Bechara, H. Damasio, D. Traneland, and A. Damasio. Deciding advantageously before knowing the advantageous strategy. *Science*, pages 1293–1295, 1997.

[Bratman *et al.*, 1988] M.E. Bratman, D.J. Isreal, and M.E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(4), 1988.

[Braubach *et al.*, 2005] L. Braubach, A. Pokahr, and W. Lamersdorf. Jadex: A BDI agent system combining middleware and reasoning. *Ch. of Software Agent-Based Applications, Platforms and Development Kits*, 2005.

[Castelfranchi and Lorini, 2003] C. Castelfranchi and E. Lorini. Cognitive anatomy and functions of expectations. In *Proc. of IJCAI03 workshop on cognitive modeling of agents and multi agent interaction*, 2003.

[Castelfranchi *et al.*, 2006] C. Castelfranchi, R. Falcone, and M. Piunti. Agents with anticipatory behaviors: To be cautious in a risky environment. In *Proc. of European Conf. on Artificial Intelligence*, Trento, Italy, 2006.

[Castelfranchi, 2005] C. Castelfranchi. Mind as an anticipatory device: for a theory of expectations. In *Proc. of the AAAI 05 Fall Symposium: From reactive to Anticipatory Cognitive Embodied Systems*, 2005.

[Kinny and Georgeff, 1991] D. Kinny and M. P. Georgeff. Commitment and effectiveness of situated agents. In *Proc. of the 12th International Joint Conf. on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, 1991.

[Kinny *et al.*, 1992] D. Kinny, M. Georgeff, and Hendler J. Experiments in optimal sensing for situated agents. In *Proc. of the 2nd Pacific Rim International Conf. on AI (PRICAI-92)*, 1992.

[Macedo and Cardoso, 2004] L. Macedo and A. Cardoso. Exploration of unknown environments with motivational agents. *3rd Int. Conf. on Autonomous Agents and Multi Agent System(AAMAS04), NYC, USA*, 2004.

[Meyer *et al.*, 1997] W. Meyer, R. Reisenzein, and A. Schutzwhol. Towards a process analysis of emotions: the case of surprise. *Motivation and Emotion*, 1997.

[Ortony and Partridge, 1987] A. Ortony and D. Partridge. Surprisingness and expectation failure: What is the difference? In *Proc. of the 10th Inter. Joint Conf. on Artificial Intelligence*, pages 106-108, Los Altos, CA, 1987.

[Parunak *et al.*, 2006] H.V.D. Parunak, R. Bisson, S. Brueckner, R. Matthews, and J. Sauater. A model of emotions for situated agents. In *Proc. Autonomous Agents and Multi-Agent Systems (AAMAS06)*, Hakodate, Japan, 2006.

[Rao and Georgeff, 1995] A.S. Rao and M.P. Georgeff. BDI agents: From theory to practice. *Proc. of the 1st conf. on MAS (ICMAS95)*, 1995.