Improving Situated Agents Adaptability Using Interruption Theory of Emotions

Clément Raïevsky and François Michaud

Department of Electrical Engineering and Computer Engineering
Université de Sherbrooke, Québec CANADA J1K 2R1

Abstract. Emotions play several important roles in the cognition of human beings and other life forms, and are therefore a legitimate inspiration to provide adaptability and autonomy to situated agents. However, there is no unified theory of emotions and many discoveries are yet to be made in the applicability of emotions to situated agents. This paper investigates the feasibility and utility of an artificial model of anger and fear based on Interruption Theory of Emotions. This model detects and highlights situations for which an agent's decision-making mechanism is no longer pertinent. These situations are detected by analyzing discrepancies between the agent's actions and its intentions, making this model independent from the agent's environment and tasks. Collective foraging simulations are used to characterize the influence of the model. Results show that the model improves the adaptability of a group of agents by simultaneously optimizing multiple performance criterion.

Introduction

In spite of significant evidence that emotion plays several crucial roles in cognitive processes [1][2][3][4], no consensus currently exists about a unified theory from which an artificial model can be derived. Therefore, to enhance our understanding of existing theories and to appreciate their usage and effects, it is still appropriate to implement them on artificial systems. Among research related to process models of emotions [5] which are applied to situated agents, we can find emotional mechanisms aimed at enhancing interaction quality between humans and synthetic agents [6], increasing synthetic agents learning abilities [7], and improving coordination among situated agents [8][9]. In these previous works, emotions are generated in two ways : by detecting specific features in the environment [6][7], or monitoring specific task progress variables without taking into account the agent's intentions [8][9]. Once generated, emotions either directly modify the agent's behavior [8][9][6], or influence other cognitive process of the agent decision-making architecture [7]. These models of emotions are limited in their versatility because they are specific either to environmental conditions for emotions generation or to mission objectives for emotional responses. However, emotions should be derived from a generic model to capture the fact that different situations can lead to the same emotions, and that the same situation can lead to different emotions. To our knowledge, no environment-independent and task-independent artificial model of emotions has yet been validated, and that is the purpose of our research.

This paper presents an artificial model of anger and fear that reproduces functions of emotions identified by the Interruption Theory of Emotions (ITE) [2][3]. This theory has not yet been implemented in situated agents. ITE states that emotions are elicited when the current decision-making process of an individual is not adapted to the experienced situation [2][3][1].Our model detects these conditions and generates emotions by monitoring temporal models of the agent's intentions. These models are independent of the agent's environment. The main function of emotions identified by ITE is to highlight the cause of the current emotional state [1][2][3]. In our model, the cause of the current emotional arousal is determined by an analysis of the agent's intentions. Identifying this cause allows the agent's motivations to adapt the agent's intentions. This signaling process is independent of the agent's tasks.

Section 1 presents ITE, followed by the description of our model in Section 2. To demonstrate and evaluate this model, it has been implemented in a behavior-based cognitive architecture and applied to a collective foraging task. Section 3 presents the experiments carried out and the results, illustrating that our emotional process improves the adaptability of a group of agents.

1 Interruption Theory of Emotions

ITE has been primarily developed by Hebb [2] and Mandler [3]. It states that interruptions of ongoing cognitive or behavioral activity trigger the *arousal* of the sympathetic nervous system which is the beginning of an emotional experience. This is also supported by the Affect Control Theory and the Self-Discrepancy Theory [10]. ITE identifies three main sources of interruption:

1. Experiencing an unexpected effect of a behavior. This occurs when conditions hinder or prevent a behavior from carrying out its function.
2. Experiencing conflicting intentions. This occurs when different decisional processes generate incompatible intentions, i.e. intentions that cannot be carried out simultaneously by actions.
3. Experiencing an unexpected situation, not anticipated by a predictive model of the world.

According to ITE, the main function of the arousal triggered by the occurrence of such interruptions is to signal to the individual that events in the environment require attention and adjustment. This generic arousal is followed by the orientation of the individual's attention toward the cause of the arousal. This highlighting process allows the individual to focus on the cause of the emotion and take the appropriate actions accordingly. Unlike other cognitive theories of emotions ITE focuses on elicitation of emotion rather than on elicitation of the different emotions. It is thus not a complete theory of emotions and a model based on this theory should be extended by models of other aspects of human emotions such as appraisal [4] and stimulus analyses to get the full range of functions associated with human emotions.

2 Artificial Model of Emotions

The artificial model of emotions we have developed aims at detecting and highlighting interruptions of cognitive or behavioral activities in order to trigger an adaptive reaction when 'normal' decision-making is no longer pertinent. This relates to ITE's first two sources of interruption. Figure 1 illustrates the hypotheses we make about the decision-making architecture of an agent to design the process which implement our model of emotions. Actions of the agent are derived from concurrent processes (i.e., Behaviors) which are activated by an Action Selection mechanism according to the agent's Intentions. The agent must have cognitive processes (i.e., Motivations) responsible for generating Intentions and determining their desirabilities. Intentions are data structures which represent particular activations and configurations of one or several Behaviors. The information regarding which Intentions are realized by the agent's actions must be available to the decision-making processes.

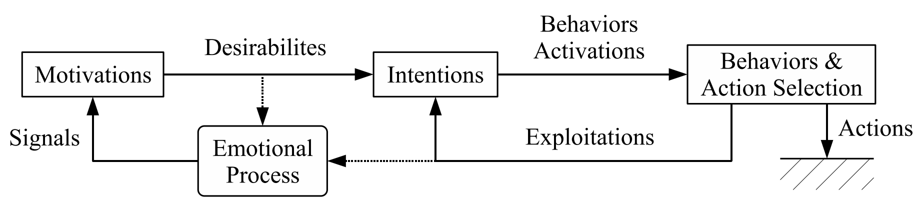

Fig. 1. Cognitive processes and concepts required by our model of emotions.

The emotional process we have developed unfolds as follows: when a situation needing adaptation occurs, an interruption is detected by the appearance of a discrepancy between the agent's intentions and the way they are satisfied by its actions (exploited). Once an interruption is detected, its cause is identified by an analysis of the current agent's intentions. The occurrence of an interruption and its cause are then signaled to the agent's motivations, which can change the agent's intentions accordingly.

Coherence between intentions and actions is checked through the monitoring of temporal models of intentions' exploitation. These models depend on the type of intentions: Goal-Oriented intentions are related to behaviors which make the agent accomplish actions aimed at fulfilling it's goals and Safety-Oriented intentions are there to keep the agent away from problematic situations. Therefore, a Goal-Oriented intention has an exploitation model of being exploited when desirable and, conversely, a Safety-Oriented intention conform to its exploitation model when not exploited.

The accumulated time $a_I(t)$ during which intention I does not conform to its exploitation model at time t is expressed by (1) and (2):

$$a_I(t) = \int_{-W_I}^{0} b_I(t)dt \tag{1}$$

$$b_I(t) = \begin{cases} 0 \text{ if intention } I \text{ conform to its exploitation model} \\ \\ 1 \text{ otherwise} \end{cases} \tag{2}$$

where W_I is the length of the sliding time window over which intention I is monitored. An interruption is detected when $a_I(t)$ becomes greater than a time threshold.

Anger is elicited by interruptions involving Goal-Oriented intentions, and conversely Fear is elicited by interruptions involving Safety-Oriented intentions. These emotions are used to modify the agent's behavior through its motivations, but the associated results are beyond the scope of this article.

Detecting and highlighting the cause of the interruption is carried out by an analysis of the agent's current intentions. It is important to distinguish the intention which *triggers* the interruption (by not conforming to its exploitation model) from the intention which is the cause of this interruption (i.e., the "responsible intention"). The responsible intention prevents the triggering intention to conform to its exploitation model and is, therefore, the subjective source of the interruption. The nature of the triggering intention determines which intention is responsible of the interruption. A Goal-Oriented intention triggers an interruption if it is desirable but not exploited during a certain period of time. Therefore, the responsible intention is the one which is being exploited the most during the recent past because it hinders the exploitation of the triggering intention. Conversely, a Safety-Oriented intention triggers an interruption because it has been exploited during a certain period of time. The responsible intention is, in this case, the triggering one. Once identified, the responsible intention is signaled to the agent's motivations as the cause of the interruption. The agent's motivations are then responsible for the adaptation of the agent behavior, keeping the model independent of particular adaptive reaction and of the agent's mission.

Interestingly, the exploitation models of intentions used by the interruption detection process are independent of the way intentions are carried out by actions; only the intentions' nature is taken into account and this is independent of the expected specific effects of the intention on the environment. Furthermore, these models are independent of the situation experienced by the agent because its perceptions are not taken into account. However, a model of emotions cannot be completely disembodied and independent from reality because emotions are not pure cognition. The emotional process we have developed is grounded in agent's reality for two reasons: first, the model parameters are time periods (i.e., a time window length and a time threshold) and capture the 'normal' operation of the agent. They therefore are bounded to the agent reality (and can then be determined by a designer through observation and measurement instead of a tedious trial and error process). Second, the emotional process originates in the intentions' exploitations which is the result of the interaction between the agent's action selection process and its environment.

The key strength of the emotional process described here is its independence from both specific effects of intentions on the environment and from reactions triggered by interruptions. This independence implies that our model is not coupled with the way the agent's behaviors carry out its intentions and ensure its applicability to other behaviors and other missions.

3 Experiments and Results

Multi-agent foraging is a widely used task with clear metrics to evaluate performance (e.g., physical interferences, traveled distance, time to complete). It is therefore suitable to illustrate our emotional mechanism. Fig. 2 illustrates the simulated environment (implemented in Stage [11]) used for the foraging experiments. The simulated agents are Pioneer 2 DXs in a pen of 6×10 meters. Six agents have to collect 12 pucks and take them one-by-one to the home region. Each agent is given two simulated sensors: one laser range finder with an 8 meter range and 180° of field of view, and one fiducial finder which returns the identifier and relative position of objects with a fiducial tag, in a range of 5 meters and a 180° field of view. Each agent has an unique fiducial identifier which allow them to perceive others' relative positions. Home flags and pucks have also fiducial ids. Agents are considered to be homogeneous, i.e., they all have the same physical and decisional capabilities. They can communicate with each other using broadcast mode (through network link).

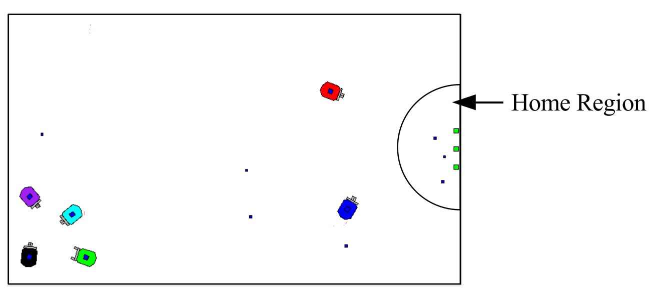

Fig. 2. Experimental setup for multi-agent foraging.

To apply our emotional process to this mission, we integrated it in a modular decision-making architecture called Motivated Behavioral Architecture (MBA) [12]. In MBA, the Behaviors are independent modules issuing commands based on the agent's perception and Intentions. Behaviors issue commands only if they are activated. Their activations and parameters are derived by a *Selection* module from the agent's Intentions. These intentions are generated by the *Motivational Modules* (MM) and are stored in the *Dynamic Task Workspace* (DTW). They are organized in a tree-like structure according to their interdependencies, from high-level/abstract intentions to primitive/behavior-related intentions. MMs are asynchronous, independent modules that can add Intentions, modify or monitor their parameters, and give activation recommendations about them. These recommendations correspond to the desirabilities of Intentions according to MMs and can take three different values: positive, negative and undetermined. The *Selection* module applies a policy to these recommendations to determine the behaviors' activations; a behavior is activated if its corresponding intention has at least one positive and no negative recommendation. The associations between Intentions and Behaviors is implemented in a *System Know-How* (SNOW) module. This module is also in charge of updating information about exploited Intentions, i.e., Intentions which are being carried out by the agent's actions.

For this mission, the five behaviors, arbitrated using subsumption, are (in order of priority):

- **Escape** makes the agent turn on itself to find a safe passage to leave the current location.
- **Obey** makes the agent execute a particular action such as stopping or turning left, according to a parameter associated with the agent's intentions.
- **Avoid** makes the agent move safely in the environment by avoiding obstacles using the laser range finder readings. Only obstacles within a 0.9 meter radius of the agent are taken into account.
- **Forage** tracks pucks, collects them one at a time and takes them back to the home region.
- **Move Forward** gives the agent a constant linear velocity.

The **motivational modules** and intentions they manipulate are:

- **Survive** ensures the security of the agent by adding and recommending the high-level *Stay Safe* intention, and specifying it by adding *Avoid* intention or *Escape* intention as its child. These two intentions are directly associated with behaviors.
- **Curiosity** makes the agent explore its environment by recommending the *Explore* intention, associated with **Move Forward**.
- **Forage** manages the foraging task by recommending the *Forage* intention (which is associated with the **Forage** behavior) and inhibiting the *Stay Safe* intention when a puck is about to be collected by the agent.
- **Social** carries out the group coordination strategy which is based on a dominance hierarchy. When an agent perceives a higher-ranked agent in a range of 1.5 meter in front of it, it stops (through the activation of the **Obey** behavior with its parameter set to stop). This distributed strategy aims at avoiding physical interference while minimizing distance traveled.

The emotional process has been implemented in a separate module, called the *Emotional Module* (EM). To detect interruptions, the EM monitors the *Stay Safe* and *Forage* intentions. The *Forage* intention has a time window length of 70 seconds and and time threshold of 60 seconds. *Stay Safe* has a time window length of 140 seconds and a time threshold of 120 seconds. These parameters have been fixed from pre-experiments trials by observation of intentions exploitations during normal situations.

These models have shown useful in detecting and adapting to two typical kinds of situation. The first typical situation occurs when the agent is in a high obstacle density area. In this situation, the *Stay Safe* intention is used often enough to trigger an interruption from the under-exploitation of the *Forage* intention or an over-exploitation of the *Stay Safe* intention. In both cases, the responsible intention is *Stay Safe* and is signaled as the cause of the interruption to the MM. The *Survive* MM uses this signal to adapt the avoidance strategy by switching the children intention of *Stay Safe* in the DTW from *Avoid* to *Escape*.

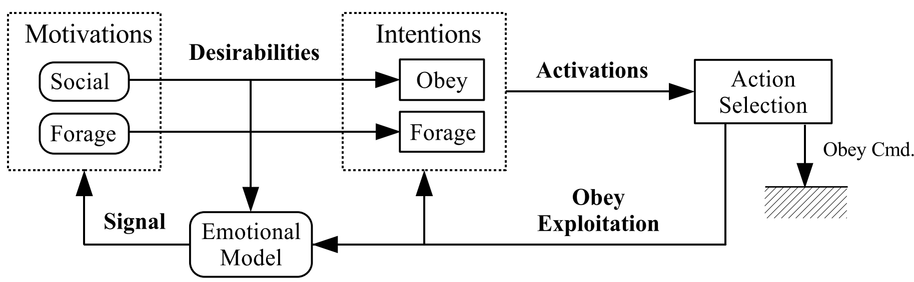

Fig. 3. Adaptation Process Example

The second typical situation occurs when an agent stops because it perceives a superior agent which is experiencing some kind of failure. Fig. 3 presents the concepts involved in the adaptation mechanism triggered by this kind of situation. This adaptation process is triggered by the prolonged exploitation of *Obey* intention which prevents the *Forage* intention from being exploited. The *Forage* intention then generates an interruption after not being exploited during 60 seconds over the last 70 seconds. Because *Forage* is a goal-oriented intention, the EM looks for the most exploited intention as the cause of the interruption and find **Obey**. This intention is then signaled to the agent's motivations as the cause of an interruption. This intention has been added by the Social MM to enforce the social rules of the group. The Social MM is therefore responsible to adapt this intention in response to the emotional signal. To do so, it triggers an update of the dominance link established between the superior and the inferior agent. The result of this update depends on the emotional state of the superior agent : if its dominant emotion is fear (meaning it is experiencing an unwanted situation involving its security) the agent is lowered in the hierarchy. Conversely, if it is experiencing anger, the agent keeps its rank in the hierarchy. This update only changes the relationship between the two involved agents.

Work by Murphy *et al.* [9] and Parker [8] are closely related to ours as they trigger behavioral adaptation from affective evaluation of task progression. However, in their work, emotion related variables are generated from an analysis of the agent's perceptions and from dedicated social messages. Their models are therefore tightly coupled with the agent's environment and with its task. Furthermore, adaptations triggered by their emotional mechanisms are dedicated to group task allocation. The main difference with our work is that our model is intended to be a generic self-analysis mechanism allowing an agent to detect problematic situations, whereas their emotional mechanism aim at improving performances of specific algorithms.

3.1 Results

To characterize the influence of our emotional mechanism while neutralizing the influence of other architectural components, we compared performances of four controllers which differ only by the presence or absence of the Social MM and of the EM: 1) the Control Group (CG) does not use the Social MM nor the EM; 2) EM refers to the controller only using the EM without the Social MM, and is used to characterize the influence of the EM without social coordination; 3) CG-So is a controller using the Social MM without the EM, and is used as a reference to evaluate the coordination strategy performance during the foraging task; 4) EM-So has both the EM and the Social MM, and is used to characterize the combined influence of the EM on the avoidance behavior and on the coordination strategy. Each controller has been used on the same series of forty randomly generated initial positions of agents and pucks to eliminate the influence of the initial conditions on the group's performance. Table 1 summarizes the observed results in terms of the following metrics. Success Rate is the ratio of failed trials (determined when pucks remain to be collected after 30 min) over the total number of trials. This metrics captures the ability of the group to recover from situations that cause it to fail and that we have not anticipated reflecting the adaptability of the group. Physical Interference Ratio is the part of time spent at a distance of 0.7 m from other agents. This metric represents the risk of collision between agents. Completion Time is the time spent to take all pucks to the home region. Traveled Distance is the total distance traveled by the agents during one trial. All these metrics, except Success Rate, take only successful experiments into account.

Comparing controllers both with the EM (i.e. EM and EM-So) and without (i.e. CG and CG-So) shows that the presence of the EM has improved the Success Rate of the group with and without coordination. This therefore suggests that our emotional process improves both the adaptability of the avoidance behavior and of the coordination strategy. As expected, the introduction of the Social MM has reduced the physical interferences between agents, making them safer, and the total distance they have traveled to complete the task, making them more efficient. However, without the EM, the coordination strategy has dramatically

Table 1. Experiments results

Control Type	CG	EM	CG-So	EM-So
Success Rate	90 %	95 %	54.8 %	87.5 %
Physical Interference Ratio	26.4 %	24.1 %	15.6 %	16.8 %
Completion Time	439	462	467	429
Traveled Distance	375	404	338	327

reduced the Success Rate of the group. This can be explained by the recurrent occurrence of an endless situation, illustrated in Fig. 2. In this situation, an agent (the black one in the bottom left corner), carrying a puck, is surrounded by inferior agents which are thus blocking it. Adding the EM to the Social MM has balanced the effects of this typical situation, bringing the Success Rate, Physical Interference Ratio, Completion Time, and Traveled Distance metrics to an optimum.

These results also show that emotions we have generated can be used to partly replicate the structuring function of emotions in some human groups [13]. It does so by keeping hierarchy relations between individuals coherent with the situation experienced by the group allowing it to adapt to the situation.

4 Conclusion

Situated agents adaptability ultimately depends on the detection of the situations for which their decision-making is not pertinent and which require a behavioral or cognitive reaction. This detection is a key problem for situated agents because their environment is dynamic, continuous and unpredictable. Psychologists have identified that one of the human emotions' functions is to highlight this kind of situation, allowing other cognitive processes to address them. We have developed an emotional mechanism which allows situated agents to detect this kind of situation by using temporal models of intentions. One of the key strength of this mechanism is that it increases the agent's adaptability without introducing either specific knowledge about the environment or about the tasks. Results from simulation experiments show that agents can recover from the malfunction of two specific algorithms (i.e., an avoidance behavior and a coordination strategy) through the use of our emotional mechanism. This has been achieved without relying either on specific knowledge about these algorithms or about specific features of the environment. The independence between our model of emotions and these algorithms guarantees the applicability of our emotional process to other algorithms and by extension to other applications. Our emotional process can be extended by adding other kind of intentions analyses such as observation of intentions resulting status, desirabilities oscillations (eliciting confusion) or lack of change in intentions (eliciting boredom) for example. We believe this versatility allows us to see our emotional process as the basis of a generic self-analysis mechanism allowing situated agents to detect and then adapt to situations for which their actions or decisions are not pertinent. Such a generic mechanism will be useful to autonomously trigger modification of decisional processes which are needed by situations or environments not anticipated by human designers and therefore, bringing artificial systems closer to complete autonomy.

5 Acknowledgments

The authors gratefully acknowledge the contribution of the Canada Research Chair, the Natural Sciences and Engineering Research Council of Canada and the Canadian Foundation for Innovation, in the support of this work.

References

1. Frijda, N.H.: The emotions. Studies in Emotion and Social Interaction. Cambridge University Press, Cambridge (1986)
2. Hebb, D.O.: The organization of behavior : a neuropsychological theory. A Wiley book in clinical psychology. Wiley, New York (1949)
3. Mandler, G.: Mind and Body: Psychology of Emotion and Stress. W. W. Norton, New York and London (1984)
4. Scherer, K., Schorr, A., Johnstone, T.: Appraisal theories of emotions: Theories, methods, research. Oxford University Press, New York (2001)
5. Scheutz, M.: How to determine the utility of emotions. In: Proceedings of AAAI Spring Symposium. (2004)
6. Velásquez, J.: A computational framework for emotion-based control. In: Proceedings of Fifth International Conference on Simulation of Adaptive Behaviors (SAB'98). (1998)
7. Gadanho, S.C.: Emotional and cognitive adaptation in real environments. In: Symposium ACE'2002 of the 16th European Meeting on Cybernetics and Systems Research, Vienna, Austria (2002)
8. Parker, L.E.: ALLIANCE: An Architecture for Fault Tolerant Multi-Robot Cooperation. IEEE Transactions on Robotics and Automation **2**(14) (1998)
9. Murphy, R., Lisetti, C., Tardif, R., Irish, L., Gage, A.: Emotion-based control of cooperating heterogeneous mobile robots. IEEE Transactions on Robotics and Autotomation **18**(5) (2002) 744–757
10. Jan E. Stets: Emotions and Sentiments. In John D. DeLamater, ed.: Handbook of Social Psychology. New York: Springer (2003)
11. Vaughan, R.T., Gerkey, B.P., Howard, A.: On device abstractions for portable, reusable robot code. In: Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems. (2003) 2421–2427
12. Michaud, F., Côté, C., Létourneau, D., Brosseau, Y., Valin, J.M., Beaudry, E., Raïevsky, C., Ponchon, A., Moisan, P., Lepage, P., Morin, Y., Gagnon, F., Giguère, P., Roux, M.A., Caron, S., Frenette, P., Kabanza, F.: Spartacus attending the 2005 AAAI Conference. Autonomous Robots, Special Issue on the AAAI Mobile Robot Competitions and Exhibition (2007)
13. Plutchik, R.: A general psychoevolutionary theory of emotion. In Plutchik, R., Kellermann, H., eds.: Emotion; Theory, Research and Experience. Volume 1. Academic Press (1980) 3–33