Adaptive agents: emotion learning

Luis Miguel Botelho

Department of Information Sciences
and Technologies
ISCTE
Lisbon, Portugal
luis@iscte.pt

Helder Coelho

Department of Computing Sciences
Faculty of Sciences
University of Lisbon
Lisbon, Portugal
hcoelho@di.fc.ul.pt

Abstract

According to [Botelho and Coelho 1998a, 1998b] emotion plays an important role at the control level of emotion systems. Essentially, emotion may lead to reflexive behaviors, to the creation or intensification of motivators, and to set new criteria for plan selection. Emotion is a sequential, possibly iterative process that comprises appraisal stages, signal generation stages and emotion-responses. In this paper, we argue that adaptive behavior may be achieved through a process of emotion learning. We point out three ways in which an agent can improve its emotion processes and therefore become more adaptive: learning new affective appraisal rules, learning new emotion-signals, and learning new emotion-responses.

1. Introduction

In this paper we discuss emotion in the perspective of the Salt & Pepper Project. The goal of the project is to develop and evaluate computer agent architectures for general intelligence. The approach taken in the Salt & Pepper Project is to build mechanisms and features that can play the same roles for artificial agents as those responsible for the success of natural agents, notably human beings. In particular we have taken inspiration in cognitive science and in neuroscience.

Emotion has been identified as one of the key elements of the intelligence and of the adaptive nature of human beings [Damásio 1994], therefore we want to develop artificial mechanisms that can play the role emotion plays in natural life. We call this mechanisms “artificial emotions”.

Since we are dealing with artificial emotions, the psychologically inclined reader should not be surprised to find out that our project addresses emotions that do not forcefully match the usual natural emotions, such as joy and relief.

The main contributions of this paper is to argue that emotion learning is a valid approach to improve the behavior of artificial autonomous agents, and to present a systematic view of the kinds of emotion learning that can take place, assuming emotion is a process involving appraisal, emotion-signal generation and emotion-response.

Emotions play important roles at the control-level of agent architectures: emotion may lead to reflexive behavior, it may intensify the motivation of an agent, it can create new motivators (e.g., new goals) and it can set new criteria for the selection of the methods and the plans the agent uses to satisfy its motives.

Since emotion is a process that operates at the level of control of an agent architecture, the behavior of the agent will improve if the emotion process of the agent improves. In this paper we explore this possibility. Namely, we explore the ways an agent can improve its emotion-processes. In section 2, we present the process view of emotion. In section 3, we relate the process view of emotion with three classes of emotion learning processes: learning appraisal structures, learning new emotion-signals, learning emotion-responses. Section 4 presents some conclusions and final remarks.

2. A process view of emotion

Within the Salt & Pepper Project, emotion is seen as a sequential, possibly iterative process that comprises appraisal stages, signal generation stages and emotion-responses. The global state of the agent is appraised, an emotion-signal may be generated and some behavior is possibly performed in response to the signal generated. This behavior may, in turn, change the agent’s global state which will be appraised again and a new signal may be generated. This process goes on and on repeatedly.

In the appraisal stage of emotion, the global state of the agent (i.e., its internal state and the external environment) is evaluated with respect to the satisfaction of its motivation structures (e.g., instincts, goals, values and attitudes). The goal of the appraisal stage is to determine if and how the global state of the agent impairs or facilitates the satisfaction of its motivators.

Although the mainstream view of emotion in artificial autonomous agents has focused on cognitive appraisal [Elliott 1994], the Salt & Pepper Project has been more concerned with affective appraisal [Botelho and Coelho

1998b]. The affective appraisal is performed by the affective engine of the architecture, whereas the cognitive appraisal is performed by the cognitive engine. The distinction between the cognitive and the affective engines has an architectural basis [Botelho and Coelho 1998b]. Interestingly, this distinction has also been confirmed by neurological data and theories [Damásio 1994] [Balkenius and Morén 1998]. For instance, the amygdala has been found to be responsible for many affective phenomena, whereas the visual cortex has been found to maintain cognitive representations of the environment.

In the Salt & Pepper Project we have been using production rules to implement the affective appraisal component of the emotion process. However, we want to stress that the use of production rules does not mean the appraisal process is a cognitive process. In this respect, production systems are only a computational tool, not a philosophical stance. The main difference between the cognitive engine and the affective engine is not the nature of the processing [Sloman 1987]. The main difference is the purpose of the processing, the kind of information that is processed, and the typical reaction time [Botelho and Coelho 1998b].

If the global state matches the left hand side of an adequate appraisal rule, an emotion signal is generated.

Some of the rules used in the affective appraisal are directly programmed into the agent, other rules may be learned in a number of ways (sections 3.1.1, 3.1.3 and 3.2) and yet some other rules result of rule compilation processes (section 3.1.2).

During the affective appraisal, explicit comparisons between the motives of the agent and its global state are avoided. Explicit comparisons involve long chains of rules which is incompatible with the demanding time constraints that are imposed by the dynamic environment on the agents that inhabit it. Instead of the long chains of appraisal rules required by explicit comparisons, the affective appraisal process tends to use single heuristic rules that implicitly determine how the global state of the agent impairs or facilitates the satisfaction of its motives.

The affective engine of the agent architecture must have access to at least a limited region of the agent episodic memory, since it must have the information regarding what has happened to it.

The features extracted by the "affective sensors" from the external environment are different and more restricted than the features extracted by the "cognitive sensors". This heuristic option enables faster processing times and therefore faster reactions.

In certain conditions, the affective appraisal stage generates signals that are intended to inform the agent that certain measures must be taken so that its global state may evolve in adaptive ways, that is, in ways that are consistent with the motives of the agent. These signals are called emotion-signals. In the Salt & Pepper Project we have been dealing with several classes of emotion-signals, such as attention-shift warnings, performance-evaluation signals, malfunctioning-component warnings, motivation intensifiers [Botelho and Coelho 1996a, 1998a, 1998b] and communication modulation signals.

Some of the emotion-signals are directly programmed into the agent, others may be learned as described in section 3.2.

When an emotion-signal is generated it is sent to the behavioral system of the agent. In certain conditions, the behavioral system reacts to the emotion-signal selecting some overt or covert behavior. Sometimes, a reflexive (overt or covert) behavior is performed. Other times, the behavioral system responds to the emotion-signal, creating a motivator that will condition the future behavior of the agent. Another way the behavioral system may react to an emotion-signal is by setting new plan-selection criteria.

For each emotion-signal there may be more than just one emotion-response, but when an emotion-signal is generated, at most one emotion-response will be selected and executed. Therefore, the exact behavior of the agent when an emotion signal is generated, depends on the pattern of activation of the behavioral network. This means the agent may behave differently in the presence of identical external situations.

This context-dependent emotion mechanism doesn't involve the symbolic processing required by the construal step in the approach of Elliott *et. al.* [Elliott 1994], it just depends on the current pattern of activation of the behavioral system of the agent.

Emotion-responses may be directly programmed into the agent, they may also be improved as discussed in section 3.3.1, and they may be learned as described in section 3.3.2.

Figure 1 summarizes the emotion process as described in this section.

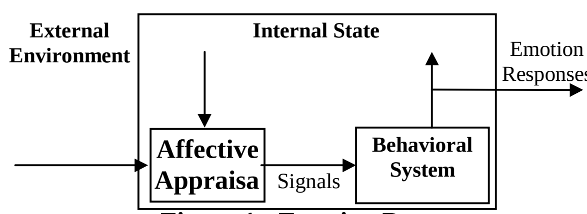

Figure 1 - Emotion Process

The cognitive loop of the agent (not fully represented in the figure) runs in parallel with the emotion loop. The main concern of the paper is the emotion loop, but some references will also be made to the cognitive component of the agent architecture.

3. Emotion learning

Emotion learning refers to the improvement of the agent's emotion process itself. Since emotion is one of the agent's control processes, improved emotion entails improved behavior.

Following the process view of emotion described in section 2, there are three main classes of emotion learning: improving and learning appraisal rules (i.e., the rules used in the appraisal stage), learning new emotion-signals, and learning or improving emotion-responses.

3.1. Improving and learning affective appraisal rules

An agent may improve its affective appraisal process in three ways. It may learn new global circumstances (external environment and internal state) in which a given emotion-signal should be generated; it may compile a long chain of appraisal rules into a single appraisal production; and it may change the characteristics of the signal that is generated in given global circumstances.

3.1.1. Learning new signal-triggering circumstances

Suppose the agent finds itself in an undesired situation. Suppose also the agent knows that what is undesirable in that situation is the fact that Q is true. Then it would have been useful if an emotion-signal were generated in a previous situation in which the agent could have decided to follow a course of action in which Q would not be true. In order for this to be possible, the agent must be able to realize that some situation is undesired; it must be able to identify what exactly is undesirable in that situation; and it must be able of determining the latest situation in which a useful emotion-signal should have been generated.

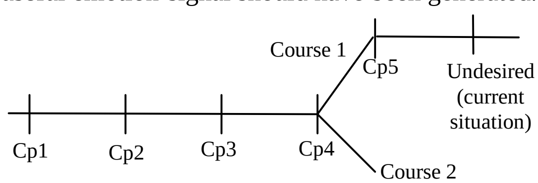

Figure 2 - An agent recent history

Figure 2 depicts a hypothetical recent history of an agent. In that history, the situations in which something could have gone different ways are called choice points. Course 1 and course 2 are two alternative courses of action. Currently, the agent is in an undesired situation. Q is what makes it undesirable. P is a specific emotion-signal.

- The current situation is undesirable because Q is the case. Possibly, Q is a conjunction of propositions, not an atomic proposition.
- In choice point cp_5, all choices available to the agent would have lead it to situations in which Q is the case (that is, the agent would not have been able to avoid Q).
- In choice point cp_4, the agent might have followed a course of action in which Q would not have been believed. However this would have been an uninformed choice, it could only have been made by chance.
- In choice point cp_3, if the agent believed P, than it would know that Q would be true in case it chooses course of action 1. Therefore, the agent should believe P in situation cp_3. Since the generation of an emotion-signal and the reaction of the agent take some time, P must be generated in situation cp_2, or in some situation before it.
- In choice point cp_2, the agent could *know* that Q would be a possibility.
- In choice point cp_1, the agent would not have means to predict that Q would be a possibility.

In summary, an useful emotion-signal (P, in the example) should have been generated in situation cp_2.

Requirements for learning new signal-triggering circumstances

1. A negative emotion-signal (R) must be generated that informs the agent that the current situation is not desirable (otherwise the agent must be able to cognitively determine that fact).
2. An attribution mechanism that determines that “Q being true” is what makes the situation undesirable. This can be made by a process that involves rationalization [Botelho and Coelho 1996b] and an explicit theory of affect, or by a process that involves deduction and an explicit theory of attribution.
3. An information processing mechanism that determines that the signal P should have been generated in situation cp_2. This mechanism must have access to the recent history of the agent’s choices, that is, the episodic memory of the agent must be available for this form of emotion learning. The agent must learn that the emotion-signal P must be generated in situations of the same class of situation cp_2.
4. When the behavioral system of the agent receives the generated emotion-signal P, it must react by enacting a behavior that avoids the course of action 1, that is, a behavior aimed at avoiding Q. In the Salt & Pepper architecture for autonomous agents, this behavior could just create the goal “avoid Q”.
5. When the emotion-signal R is generated, the agent must react (or plan to react) by initiating the attribution process described in 2 above (find Q), and by determining the triggering situation (find situation cp_2) as described in 3 above.

Question

What component of the architecture should perform processes number 2 and 3 above, the affective engine or the cognitive engine?

After the agent has learned the new emotion process described above, whenever it detects a situation of the same class of situation cp_2, its affective appraisal rules will generate the emotion-signal P. When the agent’s behavioral system receives the signal P, it may react by creating the goal “avoid Q”. Then the planning performed by the agent (whether implicitly or explicitly) will see to it that Q is avoided. The whole process is a form of adaptive behavior in which the agent learns to avoid undesirable situations.

3.1.2. Compilation of sequences of appraisal rules

This is what happens when a sequence of rules is converted in a single rule that produces the same in thhe same conditions.

When a situation is to be appraised, the affective engine searches appraisal rules whose left hand side is satisfied by the situation and whose right hand side specifies the generation of an emotion-signal. If the affective engine finds such a rule, the appraisal stops and the specified emotion-signal is generated. If the affective engine finds satisfied rules, but none of them specifies the generation of an emotion-signal, the affective engine selects one of the satisfied rules (according to some conflict resolution strategy, for instance the most specific rule), creates a new record of fired rules (to be used by the rule compilation process), records the selected rule and executes the specified action. The inference and the rule recording process proceed until an emotion-signal is generated or until no more rules are satisfied. If an emotion-signal is eventually generated by this process, the record of fired rules is handed to the rule compilation algorithm. The rule compilation algorithm generates a new rule that is asserted to the existing set of appraisal rules. The assertion mechanism guarantees that rules that generate emotion-signals are checked first by the affective system when situations are appraised.

The mechanism just described is embedded in the affective appraisal process itself, which does not require the intervention of the cognitive engine.

3.1.3. Changing the characteristics of the generated emotion-signal

The agent's affective engine may generate two different signals for the same situation if the agent learns a new emotion-signal (section 3.2) and associates it to a situation for which another emotion-signal already exists. If this is the case, it may be necessary to merge all possible emotion-signals that could be generated in the same situation into a single signal.

If it makes sense to generate both emotion-signals in the same situation, then if possible, the set of appraisal rules that give rise to both signals should be replaced by a single production that generates them both. This may happen if it is not possible to merge both signals into a single one and they both make sense.

If the emotion-signal that was learned latter is preferred to the emotion-signal that was learned earlier, then the rules that generate the signal learned earlier should be removed, and conversely if the signal learned earlier is preferred to the signal learned later.

Finally, if the agent finds out that both signals should be merged, than it must replace the appraisal rules responsible for both signals by an appraisal rule that generates the merged signal. This is possible only if the signals are of the same class but have different parameters. For instance, if one signal is an attention-shift warning with intensity I_1, and the other signal is an attention-shift warning equal to the first one except that its intensity is I_2, then the merged signal can be an attention-shift warning with an intensity given by a linear combination of the intensities of the original signals.

Requirements for learning signal merging

1. The first requirement is that the agent must be capable of preferring one signal over another one.
2. Another requirement is to determine the weights assigned to the signals, in case they should be linearly combined.
3. Both of the previous requirements can be met if the affective engine of the agent evaluates the results of generating each signal and keeps track of the evaluations. Then, if one signal is much better than the other, it would be preferred to the other. If none of them is much better than the other, they may be combined using their evaluations as their weights in their linear combination.
4. The evaluation of a signal should depend mainly on the average evaluation of the emotion-response they trigger, and also on the availability and cost of their emotion-responses.

Question

Should this process be an automatic process maintained by the affective engine, or should it involve the cognitive engine as well? The evaluation of an emotion-response may be performed by the affective engine (by means of a performance evaluation signal), the availability of the response can only be made by the affective engine, and the cost of the response may be estimated both by the affective and the cognitive engines.

3.2. Increasing the repertoire of emotion-signals

An agent may improve its emotion-process by increasing the repertoire of emotion signals that can be fired when it appraises its global state. If an agent learns a new emotion-signal, it must also learn the circumstances in which the new signal should be issued and the behaviors that might be triggered should the new signal be generated.

It does not make sense to learn a new emotion-signal for an emotion-response (overt or covert) that does not exist, but it makes sense to learn a new signal that leads to an existing emotion-response that has been found to be desirable in a given situation.

In the Salt & Pepper architecture for autonomous agents, emotion-signals are matched against a network of behaviors. Each behavior in the network is represented by a node with a header that is used in behavior selection operations. The behavior selected is the first one whose header matches the emotion-signal. Therefore, if the agent detects a situation in which it would like a certain behavior to be performed, then it must create a new emotion-signal that matches the header of the desired behavior. Then, it must create an affective appraisal rule that generates the new signal in some situation (for instance, in the situation in which the agent would like the desired behavior to be enacted).

Requirements for learning new emotion-signals

1. The agent must be able of realizing that a certain behavior (B) is desirable in a certain situation (S). This can be done in two ways. First, when the agent performs the behavior, its affective engine generates a positive performance-evaluation signal informing the agent that the behavior is desirable. Second, using hypothetical reasoning, the (cognitive engine of the) agent concludes that a certain behavior would be appropriate in situations of a certain class.
2. Create an affective appraisal rule that generates a signal that matches the header of the desired behavior (B) in situations of the class of the identified situation (S).

Questions

When should this learning process be executed? What component of the agent architecture should be responsible for it, the cognitive engine or the affective engine?

3.3. Learning emotion-responses

The agent may learn new responses to the signals that are generated by its affective appraisal system, it may improve existing behavioral responses and it may learn as a result of a response to an emotion-signal.

3.3.1. Improving existing behavioral responses

When an emotion-signal is generated and passed on to the agent's behavioral system, it is possible that a behavior is selected and executed. However, the selected behavior may reveal inappropriate for the situation at hand. In this case, the agent may learn to improve the behavior selected by means of a reinforcement learning mechanism, or it may generate a new behavior from scratch and store it in the behavioral system. This later case is analyzed in section 3.3.2. In this section, we consider the case of improving an existing emotion-response.

The main requirement for this class of emotion learning is that the agent's affective engine is capable of generating negative performance-evaluation signals whenever the enacted behavior reveals inappropriate, and positive performance-evaluation signals whenever the selected behavior performs exceedingly well (according to some criteria). In response to these performance evaluation emotion-signals, the agent should run a reinforcement learning algorithm to improve the previously selected behavior.

3.3.2. Learning new responses

When an emotion-signal is generated and sent to the agent's behavioral system, it is possible that no behavior is found that matches the signal generated (maybe not because there is no such behavior in the agent's repertoire, but because the resources assigned to find it were not enough). In this case, the emotion-signal is generated but the agent doesn't do anything because no behavior is selected.

If the agent comes to realize that it should have done something, then it can first determine what should have been achieved by an useful action (this is a goal in terms of most agent and planning architectures); and second, how this could be achieved (that is, what actions should be performed). The later is the classic planning problem, the former is a decision problem. Either the execution of the plan of action or the creation of the goal to be achieved can be suitable responses to the generated emotion-signal. The agent may chose one or the other according to the features of the situation and depending on whether or not it has actually built the plan of action the time it realizes that something should have been done. Goal creation behaviors are more flexible strategies, but result in slower overt reaction (it requires that some kind of planning be done on the fly). Reflexive action is less flexible but is faster.

Once the agent has determined the response to the generated emotion-signal, it must add this new behavior to its repertoire and associate it to the emotion-signal that was generated. In terms of the Salt & Pepper architecture for autonomous agents, this means to create a new node with this new behavior and store it in its behavioral network. The header of this new node must match the emotion-signal that was generated and left unanswered.

Requirements for learning new emotion-responses

1. The agent must be able of recognizing that an emotion-signal was generated and no behavioral response was produced. This is an easy to solve architecture dependent problem.
2. The agent must be able of realizing that something should have been done in response to the generated emotion-signal. As usual, this can be achieved in two ways. First, the fact that the generated emotion-signal was not "attended" may be noticed by the affective engine and another signal may be generated. Second, the agent may cognitively recognize this need.
3. The agent must have some decision making and some kind of planning capabilities.

3.3.3. Learning as a result of a response to an emotion-signal

When the emotion-response of the agent to some emotion-signal is an adaptive learning process, the agent improves its behavior as a result of executing the response to the emotion-signal. Therefore this is not a case of emotion learning, thus it will not be considered in detail in this paper. This is a case in which emotion leads the agent to improve its behavior. In principle, any learning algorithm can be triggered as a response to an emotion-signal. Just to mention a few possibilities, the agent may trigger a rationalization process that explains why some event has happened and learn to avoid the causes of the undesirable situation (or to try to achieve them, in case of a desirable situation), it may also mark some decision or action selection rules as "bad rules in some context", and it may set new criteria for plan selection that avoids interruptions (and hence, avoid experiencing new negative emotions). It is worth noting that this later form of learning corresponds to the proposal of Herbert Simon presented in [Simon 1967].

4. Conclusions and final remarks

Our work on artificial emotion suggests some answers to some questions regarding autonomy, emotion and cognition interplay, and affective systems synthesis.

First, it shows that the creation of new motivators (e.g., goals) may result from the emotion process (actually, the emotion response). This is part of the answer to a repeatedly asked question by the autonomous agents community: "where do goals come from?".

Second, it reveals that at least part of the episodic memory of the agent must be accessible to emotion and to emotion learning.

Finally, it puts forth some principles for the design of the affective appraisal system of the agent: avoid explicit comparisons between the global state and the agent's motivators, replace long chains of rules by single appraisal productions, restrict the appraisal of the global state to a limited amount of features of the internal state and the external environment.

In this paper, we argue that emotion learning may improve the agent behavior, since emotion plays fundamental roles at the level of control of the agent architecture. This amounts to saying that better control entails more adaptive behavior.

The main contribution of the paper is the description of seven processes by which an agent may learn emotion. These processes belong to three major classes: learning affective appraisal rules, learning new emotion-signals, and learning emotion responses. These classes of emotion learning are derived from the model of emotion presented in [Botelho and Coelho 1998a, 1998b]. This model view emotion as a process with three main components: appraisal, emotion-signal generation and emotion-response.

In most cases, the main architectural requirement for an emotion learning process to be possible is that situations, events and behaviors are detected and evaluated, which is one of the roles of the emotion process itself. In some cases however, the cognitive engine is also required for emotion learning. This is what may happen, for instance, when the agent has to perform attributions (section 3.1.1), make decisions (section 3.3.2), or figure out plans of action (section 3.3.2).

The systematic analyses presented in the paper may be used as a guide for developing more adaptive agent architectures, based on emotion learning.

Acknowledgments. This work was partially supported by the contracts PSCH/OGE/1038/95 (Mago), 2/2.1/TIT/1662/95 (Sara) and by UNIDE/ISCTE.

5. References

Balkenius, C. and Morén, J. (1998) "A computational model of emotional conditioning in the brain", in *SAB'98 Workshop "Grounding Emotions in Adaptive Systems"*, Forthcoming.

Botelho, L.M. and Coelho, H. (1996) "Emotion-based attention shift in autonomous agents", in Müller, J.P., Wooldridge, M.J. and Jennings, N.R. (eds) *"Intelligent agents III. Agent theories, architectures and languages" (ATAL96)*, p277-291, Springer-Verlag

Botelho, L.M. and Coelho, H. (1996) "Agents that rationalize their decisions", *ICMAS'96*, p3-10

Botelho, L.M. and Coelho, H. (1998) "Artificial autonomous agents with artificial emotions", *Proc. of the Second International Conference on Autonomous Agents (Agents'98)*, p449-450.

Botelho, L.M. and Coelho, H. (1998) "Artificial emotions", *Department of Information Sciences and Technologies, ISCTE, Lisbon, Portugal, Technical Report ISCTE-DCTI-1998-002*

Damásio, A.R. (1994) "Descartes' Error: Emotion, Reason and Human Brain", Picador, London

[Elliott, C. (1994) "Research problems in the use of a shallow artificial intelligence model of personality and emotion", *AAAI'94*, p9-15

Simon, H.A. (1967) "Motivational and emotional controls of cognition", *Psychological Review*, 74:29-39

Sloman, A. (1987) Sloman, A. (1987) "Motives, mechanisms and emotions", *Cognition and Emotion*, 1:217-234