Salt & Pepper Architecture and Toolkit

Bruno Gonçalves (bruno.goncalves@iscte.pt)
Nuno Jesus (nuno.jesus@iscte.pt)
Luis Botelho (luis.botelho@iscte.pt)

We, the Body and the Mind Research Lab
ADETTI/ISCTE
Av. das Forças Armadas, Edifício ISCTE
1600 Lisboa, Portugal

Abstract

The main goal of this paper is to describe a set of computational and conceptual tools that were developed and used to create autonomous artificial agents using the Salt & Pepper architecture, which is a cognitive architecture inspired in descriptive models of decision making and judgment from cognitive science and cognitive social psychology, and in recent theories and findings from neuroscience.

Introduction

Few cognitive and neurological theories have been used as the source of inspiration for the development of artificial autonomous agents. SOAR [Newell 1994][Rosenbloom 1993] and ACT-R [Anderson et al 2002][Anderson and Lebiere 1998] are noticeable exceptions. However SOAR has long departed from its cognitive roots and ACT-R is not exactly an agent architecture. Although ACT-R can be used as an agent architecture, its main purpose is the study of human cognition, not to create artificial intelligent agents.

Salt & Pepper [Botelho and Coelho 2001] is a cognitive architecture with some similarities with ACT-R. However, it is intended as an agent architecture not as a model of human cognition. The purpose is to create effective and efficient agents with general intelligence.

Besides having a different purpose, Salt & Pepper has another important difference with respect to ACT-R: it is as concerned with cognition as with emotion, while ACT-R is more concerned with cognition (although some papers have been written on this topic by the ACT-R community [Belavkin 2001][Ritter, Belavkin, and Elliman 1999]).

Architecture	Object	Cognition	Emotion
ACT-R	Human cognition	Yes	No
SOAR	Agent architecture	Yes	No
S & P	Agent architecture	Yes	Yes

Figure 1 –Salt & Pepper, SOAR and ACT-R

Salt & Pepper was implemented first as a simple question-answering prototype in order to study its properties and compare them with certain biases in human judgment and decision-making. The results are summarised in [Botelho 1997] and in [Botelho and Coelho 2001].

Recently, in the European IST Project SAFIRA, Salt & Pepper was re-implemented and used as an action-control architecture for agents with emotions.

The results gathered during these two stages confirmed that Salt & Pepper agents exhibit adaptive behaviour and therefore could be used in challenging environments facing demanding problems. Hence, we have developed a set of conceptual and computational tools used to help developers creating Salt & Pepper agents. In the course of developing those tools and the associated experiments, the Salt & Pepper architecture has undergone several modifications.

This paper describes the current state of the Salt & Pepper architecture and tools for artificial intelligent autonomous agents.

The next section presents an explanation of the Salt & Pepper conceptual architecture for artificial autonomous agents. The third section describes the set of conceptual and computational tools for developing Salt & Pepper agents. The fourth section reports some experiments with the set of tools created to develop Salt & Pepper agents. The fifth section discusses related work. Finally, the last section concludes and addresses topics for further research.

Salt & Pepper Conceptual Architecture

Salt & Pepper Autonomous Agents exhibit a somehow unpredictable behaviour in which deliberative action control processes are occasionally overridden by uncontrolled (involuntary) emotion-responses. Emotion is a sequential, possibly iterative process comprising the evaluation of emotion eliciting conditions, the generation of emotion signals, and emotion responses [Botelho and Coelho 2001] as showed in Figure 2.

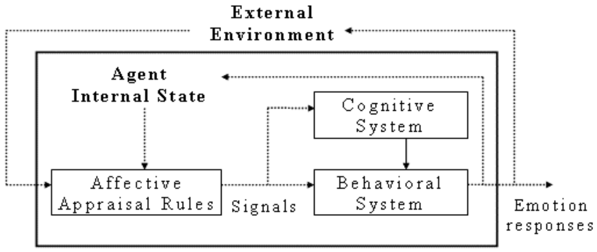

Figure 2 - Salt & Pepper emotional model

Emotion-signals are generated by emotion-eliciting mechanisms. Some of these are very simple non-symbolic processes, while some others are more elaborated cognitive appraisal processes, as hypothesized by Fridja [Fridja 1993].

Emotion eliciting conditions and the generation of emotion-signals are implemented by a set of production rules in which the firing frequency is controlled by explicit parameters dependent on the specific emotion category being represented in each rule.

The behaviour of Salt & Pepper Autonomous Agents depends on their past interactions, which provides the basis for adaptation. Adaptation is further improved by an internal punishment and rewarding process based on emotion.

Interestingly, Salt & Pepper provides a new software paradigm in which the programming of the deliberative action control mechanisms is totally separated from the programming of emotional and other non-deliberative behaviours.

Agent mind major blocks

The Salt & Pepper agent's mind is composed by some memory systems including an active long-term memory, an active working memory, input and output buffers; several emotion eliciting mechanisms; and an interruption mechanism.

The adaptive nature of the Salt & Pepper architecture is rooted in an associative long-term memory with spreading activation exhibiting the recency and the frequency effects (more recently and more frequently stimulated information is more easily found in long-term memory) [Botelho 1997][Botelho and Coelho 2001]. These context-dependent adaptive effects depend (i) on the patterns of activation of and associations between long-term memory knowledge structures, which reflect the past interactions of the agent, and (ii) on the rewards or punishments due to positive or negative emotion signals of performance evaluation.

The information unit of the Salt & Pepper long-term memory is called a node and may contain any kind of knowledge, including episodic memories, general knowledge (e.g., rules), emotion-responses, and procedures.

Working memory stores knowledge and information structures being processed at any moment. Some of the contents of working memory are recruited from long-term memory (nodes), others are read from the input buffers, and yet others are temporary cognitive structures generated by information processing mechanisms. The temporary cognitive structures are automatically removed after some time has elapsed since their creation.

The interaction between the agent mind and the sensory/effectory functions is mediated by the input and output buffers. The input buffer stores the information sent by the sensors, which can only be accessed by processes in working memory. Although the information in the input buffer is not automatically read, it is presented to long-term memory increasing the activation of nodes matching it. The output buffer is responsible for the storage of commands to be sent to effectors.

The agent's mind also contains a number of emotion eliciting mechanisms. These mechanisms evaluate received information and possibly generate emotion signals that are presented to long-term memory possibly causing activation of matching nodes to increase. In our experiments we have used emotion eliciting processes coupled with effectory mechanisms. These emotion eliciting processes, which reflect action execution status, have theoretical relevance since they cannot be considered cognitive appraisal processes.

When the execution of an action fails/succeeds, the emotion eliciting mechanism coupled with the effector generates a negative/positive emotion signal of the category *performance-evaluation*. The intensity of the generated emotion-signal depends on the number of failures and successes of the same action. Therefore the behavioural control component will receive performance evaluation emotion-signals of different intensities – sometimes the intensity will be enough to trigger an emotion response. Whether or not the failure or success to execute an action leads to an emotion-response depends of the agent past interaction with the external world.

The interruption control mechanism is responsible for interrupting the deliberative processing whenever the activation of any long-term memory node becomes high enough (e.g.., more than a certain proportion of working memory activation). When this happens, if interruptions are enabled, the contents of the interrupting node may be copied to and processed in working memory. The contents of nodes interrupting working memory are instantiated with the information that caused the interrupting node to become more activated (i.e., external stimuli or emotion signals).

If the interruption is caused by an external stimulus, it actually interrupts current working memory processing only if the responsible stimulus is still present in the input buffer at the time the interruption is considered.

If the interruption is caused by an emotion-signal, it actually interrupts current working memory processing, if the responsible emotion signal has not been received long time ago, at the time the interruption is considered.

From a software engineering point of view, the Salt & Pepper behaviour control module is a layered system, which facilitates its extensibility and use as a software development tool. This system is responsible for the deliberative control of the agent at the service of its goals. However, since the behavioural control module is implemented on top of the Salt & Pepper long-term memory and may receive emotion-signals from the existing emotion eliciting processes, the control of the agent will not be purely deliberative. Due to the presence of external stimuli or to the reception of emotion-signals, it is possible that some action or process not voluntarily selected by the deliberative control mechanism is automatically executed overriding the deliberative control.

The Salt & Pepper behavioural control module is composed of five layers (Figure 3): Memory, Cognitive Engine, Communication, Action Control and finally Reasoning Mechanism.

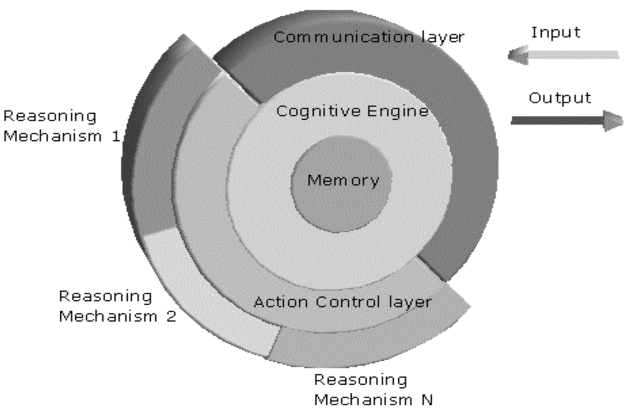

Figure 3 - Salt & Pepper layered architecture

The Memory layer consists of the Salt & Pepper long-term memory model, already described.

The Cognitive Engine includes working memory, input and output buffers, the interruption control mechanism, and interacts with long-term memory.

The Action Control Layer, which subsumes the Cognitive Engine, provides general action control functionalities to be used by any specific behavioural control mechanism, such as the emotional production system described later. Action control mechanisms are responsible for selecting the actions (internal or external) that should be executed taking into account the current motives of the agent, the external environment, and the generated emotion-signals.

The Salt & Pepper Cognitive architecture supports the co-existence of multiple action control mechanisms, some of which are less thoughtful but faster while others are slower but more deliberate.

Emotional Production System

Em-PSys, the emotional production system, is an action control module defined and implemented on top of the Action Control Layer and hence it subsumes all of its characteristics.

Em-PSys is a general-purpose emotional production system that takes a set of files containing production rules, pattern-action clauses and action descriptions for the application domain. Pattern-action clauses consist of pairs of a pattern and an action. When the pattern matches an external stimulus or an emotion-signal, the action is executed. Unlike production rules, no inference is used to match the pattern with received external stimuli or emotion-signals. Pattern-action clauses are mostly used to specify the behavioural response to generated emotion-signals.

Usually, the Em-PSys works as a regular production system. However, when an emotion signal is generated, it is likely that the pattern of a pattern-action clause matches the generated signal. If this happens, the current working memory processing can be suspended, and the action part of the pattern-action clause can be executed. In those circumstances, the deliberative behavioural control is overridden by uncontrolled emotion responses. The agent designer does not have any control over the circumstances in which the deliberative control is overridden by emotion.

Using action descriptions, the Em-PSys automatically computes the desired expected results of the actions deliberatively selected for execution. The desired expectations can be sent to emotion eliciting processes so that performance evaluation signals can be generated accordingly.

Tools for Agent Development

During the European IST Project SAFIRA, we have developed a set of conceptual and computational tools for the development of Salt & Pepper agents. The tools comprise (i) the SAFIRA Integration Framework; (ii) the SAFIRA Development Toolkit; and (iii) the SAFIRA run-time support system composed by the CRS (Central Registry Service) and by the Interaction Manager.

The SAFIRA Integration Framework adopts a component-based approach to agent development and prescribes a set of principles guiding the interaction between the components that make up an agent.

The SAFIRA Development Toolkit is an extensible graphical development environment and a set of built-in general-purpose components used to help agent designers developing their Salt & Pepper agents.

The Central Registry Service (CRS) is a component that implements a yellow pages service allowing the flexible and dynamic interaction between the components that make up the agent.

The Interaction Manager is a software module that provides run-time support at the level of the interaction between components (TCP/IP message passing, message parsing and generation, and service registration and discovery). The Interaction Manager is integrated in every component created using the SAFIRA Development Toolkit.

SAFIRA Integration Framework

The SAFIRA integration framework defines a common component interface, a set of principles that govern the general shape of the interaction between components, and an interaction architecture that supports the defined principles.

In order to allow developers to create efficient agents, Salt & Pepper has been implemented as a component-based architecture, which defines four kinds of functions for the agent components: effectory, sensory, emotion eliciting and behavioural control. Each of these functions may be implemented by one or more components; and each component may implement more than one of these functions.

The Integration Framework defines three principles that guide the integration of the components that make up an agent:

With the exception of the CRS, no component can receive information, commands or emotion-signals that were not previously requested (though a question or a subscription). Information can only be sent in reply to a question or when it is subscribed. Commands and emotion-signals can only be subscribed.

Communication uses TCP sockets and textual XML (eXtensible Markup Language) messages.

In order to keep the traffic of messages between components within an acceptable level, information received from the real world by the sensors is sent to components that have requested it only if some change has occurred since the last time the information was sent. This way, the sensors never send consecutive messages to the same component containing exactly the same information. However stimuli present in the input buffer of the Salt & Pepper behavioural control components continues to activate long-term memory nodes that match them. This way, although the information is not repeatedly sent to the behavioural control component, it produces the same memory effect as if it has been repeatedly received.

SAFIRA Development Toolkit

The SAFIRA toolkit is a software tool with graphical user interface that guides the development of complete applications using Salt & Pepper components and the integration framework principles.

The Toolkit facilitates the development of components and component-based agents with respect to aspects of components interaction, including communication (sockets, message parsing and generation), and management of the protocols used in the interaction. Significant parts of the code related to these aspects of the interaction are automatically generated by the Toolkit.

It generates the code that registers the services provided by the component in the CRS; it generates the code used to request the identities of providers of desired information, emotion signals and commands from the CRS; it generates the code used by the component to request information, emotion-signals and commands from provider components; finally it generates the headers of the methods that handle the received messages.

Most often, the effectory and sensory functions are domain dependent therefore agent designers have to implement them from scratch. The emotion eliciting and control functions may be built on top of general mechanisms such as rule-based systems, planning algorithms and case-based reasoning system, therefore the SAFIRA Development Toolkit provides general purpose built-in emotion eliciting and behavioural components, which have been developed on top of the Salt & Pepper architecture. In order to create domain dependent emotion eliciting and behavioural control components, agent designers can edit the contents of these general-purpose built-in components, using the graphical development shell of the SAFIRA Development Toolkit.

Central Registry Service

As already mentioned, the CRS (Central Registry Service) ensures run time support for deployed agents providing a yellow pages service. Each component making up an agent can register its services in the CRS. When a component needs a certain service, it consults the CRS to discover the components providing the desired service. In reply, the CRS sends a message to the requester containing the provider port and IP address. Then the requester directly contacts the provider and asks it for the desired service or information.

The services are typically requests and information subscriptions.

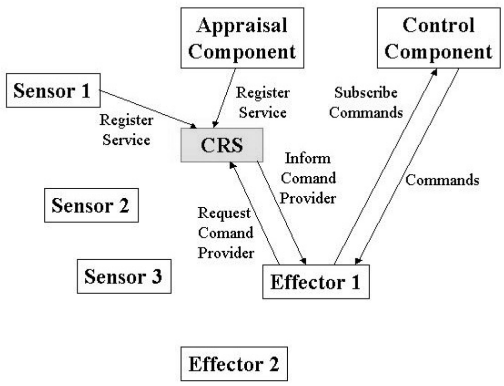

Figure 4 - Central Registry Service

The CRS allows the agent composition to dynamically change in run-time without the need to re-compile the agent.

Experiments / Applications

Salt & Pepper was first used in a prototype of a question-answering system implemented to study the relation between its properties and some of the findings about several biases observed in human behaviour. The results, summarized in [Botelho 1997] and [Botelho and Coelho 2001], shown that Salt & Pepper exhibits several human-like behaviours. Information in long-term memory can be primed by previous stimuli resulting in biased behaviour when answering subsequent questions. If a certain word (the prime) is more strongly associated to a certain piece of information B than to another piece of information C, a subsequent question that could have B and C as possible answers will most likely elicit the answer B, instead of C.

From 1998 to 2000, in the scope of the European IST Project SAFIRA, the Salt & Pepper architecture was repeatedly implemented (first in Prolog and later in Java and Prolog), and undergone important improvements. In SAFIRA, Salt & Pepper was used as an action-control architecture for artificial agents with human like characteristics as emotion. Salt & Pepper was used to implement an agent to play the Monga World game. Monga World is a demanding game conceived as a test bed used in the evaluation of Salt & Pepper both as a development paradigm and as an agent conceptual architecture. It was shown that Salt & Pepper allows a kind of programming paradigm in which the programming of deliberative control is totally separated from the programming of the non-deliberative emotional behaviour. The experience with Salt & Pepper is not enough to allow any definitive conclusions about its value as an action control mechanism.

We are planning to use the Salt & Pepper architecture to build resource limited, situation-aware agents for dynamic environments in the scope of a future European Project submitted to the sixth framework programme. First, a dynamic model of motivation strength will be built into the Salt & Pepper architecture as a dynamic mechanism to automatically allocate computation resources to each of the tasks faced by the agent. Salt & Pepper long term associative memory with spreading activation similar to those described in [Anderson and Pirolli 1984][Collins and Loftus 1975] will be used to represent the relationships between attributes, classes, ontologies and information providers that will allow an information broker agent to efficiently select information providers for received complex information requests.

Related Work

A few architectures inspired by cognitive science and by neuroscience have emerged from AI research, becoming in some cases cognitive-modelling tools with a large user community like SOAR and ACT-R.

Salt & Pepper has several similarities with such architectures.

TABASCO [Staller and Petta 1998] presents important similarities with Salt & Pepper, namely the existence of an action monitoring process evaluating its results. Another import issue is the fact that, in TABASCO, there is no direct link between emotion and action; just an action tendency. In Salt & Pepper, there is also no direct link between emotion and behavioural response. The result of the appraisal stage is an emotional signal that may or may not lead to an action, although it always changes accessibility patterns in long-term memory.

Both ACT-R and Salt & Pepper have a long-term memory with activation and associative strengths. However, in the case of ACT-R, long-term memory stores only declarative knowledge; procedural knowledge is stored in a flat memory. Salt & Pepper LTM can store any kind of knowledge.

SOAR uses a flat memory to store all kinds of knowledge therefore it is not an adaptive memory; memories do not become more or less accessible as a result of past interaction. In SOAR, adaptation is achieved mostly through a chunking process [Laird, Rosenbloom and Newell 1986].

An important characteristic distinguishing Salt & Pepper from ACT-R and SOAR is the possibility of extending it by developing new information processing mechanisms able to process and manage any kind of procedural and declarative knowledge structures. This characteristic ensures more flexibility in problem solving.

Humanoid robot projects are also a great contribution for agent architectures. The ISAC humanoid architecture [Peters et al 2001] uses a long-term memory with the same principles of Salt & Pepper and ACT-R memories. KISMET humanoid robot architecture [Breazeal 1998] has a motivational system used to influence how and when the humanoid acts to satisfy its needs.

Conclusions and future developments

This paper describes a set computational and conceptual tools developed for creating artificial intelligent autonomous agents. The main contribution of this paper is to present an alternative to cognitive-modelling tools like SOAR or ACT-R, namely by supporting an emotional theory.

Future steps: development of a variety new mechanisms, including an Emotional Case Based

Reasoning mechanism; an emotion-learning algorithm used to learn new behavioural control structures and new emotion eliciting structures; emotional interleaved planning and execution system; dynamic model of motivation to control the amount of effort to be allocated to each task faced by the agent; dynamic model of emotion (how does emotion intensity change?).

Acknowledgments

This work is partly supported UNIDE/ISCTE.
The authors wish to express their gratitude to Pedro Figueiredo and Xico Ventoinha for being so bad.

References

[Anderson and Lebiere 1998] Anderson, J. R.; Lebiere, C.J. 1998. “Atomic Components of Thought”. Lawrence Erlbaum Associates

[Anderson and Pirolli 1984] Anderson, J.R. e Pirolli P.L. (1984) "Spread of activation", *Journal of Experimental Psychology: Learning, Memory and Cognition,* 10:791-798

[Anderson et al 2002] Anderson, J. R.; Bothell, D.; Byrne M. D.; and Lebiere, C. (submitted). “An Integrated Theory of the Mind”. Psychological Review

[Belavkin 2001]Belavkin, R. V. 2001. “The Role of Emotion in Problem Solving”. In *Proc. of the AISB'01 Symposium on Emotion, Cognition and Affective Computing* (pp. 49--57)

[Botelho 1997] Botelho, L.M. 1997. "Realização de agentes inteligentes: uma abordagem baseada em modelos cognitivos de tomada de decisão" [Building intelligent agents: an approach based on the cognitive models of decision making] (in Portuguese), Ph.D. dissertation, Department of Information Sciences and Technologies, ISCTE

[Botelho and Coelho 2001] Botelho, L.M.; and Coelho, H. 2001. "Machinery for Artificial Emotions". *Cybernetics and Systems*, 32 (5): 465-506

[Breazeal 1998] Breazeal, C. 1998. “A motivational system for regulating human–robot interaction”, *Proceedings of the National Conference on Artificial Intelligence*, Madison, WI, 1998, pp. 54–61

[Collins and Loftus 1975] Collins, A.M. and Loftus, E.F. 1975. "A spreading-activation theory of semantic processing". *Psychological Review*, 82:407-428

[Frijda 1993] Frijda, N.H. (1993) “The place of appraisal in emotion”. *Cognition and Emotion*, 7(3/4):357-387

[Laird, Rosenbloom and Newell 1986] Laird, J.E.; Rosenbloom, P.S.; and Newell, A. (1986) “Chunking in Soar: The Anatomy of a General Learning Mechanism”. Machine Learning 1:11-46

[Newell 1994] Newell, A. 1994. “Unified Theories of Cognition”. Harvard University Press

[Peters et al 2001] Peters II, R. A.; Kawamura, K. .; Wilkes, D.M.; Hambuchen, K.A.; Rogers, T. E.; and Alford, A. 2001. “ISAC humanoid: an architecture for learning and emotion”, *Proceedings of the 2001 IEEE-RAS International Conference on Humanoid Robots*, pp. 451-459

[Ritter, Belavkin, and Elliman 1999]Ritter, F.E.; Belavkin, R.V.; and Elliman, D.G. 1999. “Affective Computing: The role of emotion in human-computer interaction”. In A. Monk, A. Sasse, & A.Crerar (Eds.) *Proc. of the British HCI Group one-day meeting in conjunction with University College London*

[Rosenbloom 1993] Rosenbloom, P.S. (Editor). 1993. “The Soar Papers: Research on Integrated Intelligence”. The MIT Press

[Staller and Petta 1998] Staller, A.; and Petta, P. 1998. Towards a tractable appraisal-based architecture for situated cognizers. In Cañamero, D.; Numaoka, C.; and Petta, P., eds., *Workshop: Grounding Emotions in Adaptive Systems Int. Conf. of Simulation of Adaptive Behaviour, from Animals to Animats, SAB’98*. 56–61.