Biologically Inspired Cognitive Architectures 18 (2016) 33–40

Contents lists available at ScienceDirect

Biologically Inspired Cognitive Architectures

journal homepage: www.elsevier.com/locate/bica

Research article

A cognitive-affective architecture for ECAs

Joaquín Pérez [a,*], Eva Cerezo [a], Francisco J. Serón [a], Luis-Felipe Rodríguez [b]

[a] Aragón Institute of Engineering Research (I3A), Department of Computer Science and Systems Engineering (DIIS), University of Zaragoza, Spain
[b] Department of Computer Science and Design, Instituto Tecnológico de Sonora, México, Mexico

ARTICLE INFO

Article history:
Received 8 June 2016
Revised 30 September 2016
Accepted 2 October 2016

Keywords:
Affective computing
Cognitive architecture
Conversational agent

ABSTRACT

The development of *Embodied Conversational Agents* (ECAs) involves a large number of challenges such as the modeling of cognitive and affective functions in order to achieve realism and believability in this type of intelligent agents. An approach to provide ECAs with capabilities for cognitive processing such as learning, decision making, planning, and perception has been the use of cognitive architectures. Moreover, the literature reports several affective models for the generation, classification, and management of emotions in ECAs. Nevertheless, there is a need of cognitive-affective architectures that address the problem of achieving natural interaction and realistic behavior in ECAs. In this paper, we discuss the state of the art on existing cognitive architectures, affective models, and ECAs, and propose a cognitive-affective architecture based on Soar and extended with an affective model inspired by ALMA. The proposed cognitive-affective architecture is designed to allow ECAs to include and take advantage of features such as reinforcement learning, episodic memory, and emotion management.

1. Introduction

The concepts of intelligent agent and embodied conversational agent (ECA), introduced by Cassell and Bickmore (2000), have emerged and progressed during the last years. In particular, an ECA is a virtual character with conversational and learning skills, realistic body expression, and the ability to perceive an environment, reason about it, and act accordingly. The incorporation of these characteristics in ECAs involves a variety of challenges, including physical appearance issues, realistic animations and expressions, cognitive modeling, abilities to interact with humans and other agents, and ethical issues (Kasap & Magnenat-Thalmann, 2007).

An approach to supporting all these aspects of ECAs is the use of cognitive architectures. A cognitive architecture can be defined as a scheme or pattern for structuring the functional elements that make an intelligent agent as a whole (Langley, Laird, & Rogers, 2009), so it simplifies the design of an integrated system in which all desired features can be included. In this context, cognitive architectures seem appropriate means to provide ECAs with capabilities for cognitive processing such as learning, decision making, planning, and perception.

There is nowadays a consensus among researchers in fields such as psychology and neuroscience about the close interaction between cognition and emotion in humans (Lane, Nadel, Allen, & Kaszniak, 2000; Phelps, 2006). Particularly, evidence shows that emotions play an essential role in cognitive functions such as decision making, learning, and planning (Damasio, 1994). In the field of computer science, these types of findings have led to the concept of *affective computing*, introduced by Picard (1997), which refers to any type of computing related to emotions or other affective phenomena. In recent years, this research area has reported numerous contributions (Scherer, Bänziger, & Roesch, 2010; Reisenzein et al., 2013). In this context, a crucial requirement to achieve human-like behavior in computational agents is to consider affective aspects. (Scheutz, 2004) suggests that emotions are crucial for the agent's action selection, adaptation, social regulation, sensory integration, motivation, goal management, memory control, learning, and strategic processing.

Although the literature reports several models for the generation, classification, and management of emotions (Scherer et al., 2010; Rodríguez & Ramos, 2015), these models usually disregard the importance and implications of modeling the interaction between cognitive and affective aspects for the development of computational agents with human-like behavior. Moreover, in spite of the great advances achieved in the last years in the development of ECAs, there is a lack of architectures that combine *cognitive* and *emotional aspects*. In this sense, there is a need of cognitive-affective architectures that address the problem of

* Corresponding author.
E-mail addresses: joaquinperezmarco@gmail.com (J. Pérez), ecerezo@unizar.es (E. Cerezo), fjseron@unizar.es (F.J. Serón), luis.rodriguez@itson.edu.mx (L.-F. Rodríguez).

http://dx.doi.org/10.1016/j.bica.2016.10.002
2212-683X/

achieving natural interaction and realistic behavior in ECAs. This type of cognitive-affective architecture must be designed to provide mechanisms that model the interaction between cognitive and affective processes.

In this paper, we propose a novel cognitive-affective architecture for ECAs that integrates an *emotion model* based in ALMA (Gebhard, 2005) and the Soar cognitive architecture (Laird, 2012). The proposed architecture is designed to allow ECAs to include and take advantage of mechanisms such as reinforcement learning, episodic memory, and emotion management. Importantly, this cognitive-affective architecture can be easily included in existing ECAs, due to a modular design that does not force an agent to use unwanted or unnecessary characteristics.

The paper is structured as follows. In Section 2, we present a discussion of the state of the art regarding ECAs, cognitive architectures, and affective models. In Section 3, we present the proposed cognitive-affective architecture for ECAs. Finally, conclusions are presented in Section 4.

2. Related work

In this section, we present related work about cognitive architectures, affective models, and ECAs. In particular, we discuss design aspects of representative cognitive architectures, explain how emotions have been studied and modeled in the fields of psychology and computer science, and analyze how these architectures and models have influenced the development of ECAs. The analysis of previous research in these areas provides an interesting panorama of key characteristics and aspects for the development of novel cognitive-affective architectures.

2.1. Cognitive architectures

Cognitive architectures are designed to build intelligent agents that can solve a wide range of problems using heterogeneous knowledge (Langley et al., 2009). This type of cognitive architectures usually include different specialized memories useful for the development of ECAs, such as procedural memory (defines the actions that can be performed), semantic memory (stores known information about the environment), episodic memory (stores memories of past experiences) and perceptual memory (serves to recognize and classify objects in a structured way). They also tend to include control and process components, learning mechanisms, data representations, and input/output systems. In this section, we analyze those essential aspects of cognitive architectures (e.g., memory systems, learning mechanisms, and affective aspects) based on previous work of the BICA Society (Biologically Inspired Cognitive Architectures).[1]

4CAPS (Cortical Capacity-Constrained Concurrent Activation-based Production System): is a cognitive architecture successor of CAPS and 3CAPS. The CAPS architecture was capable of building models for language comprehension and problem solving (Thibadeau, Just, & Carpenter, 1982). This architecture was replaced by 3CAPS, which added constraints in terms of the resources allowed for the system. The 4CAPS cognitive architecture extends 3CAPS by adding neuroimage measures, enabling the study of cognitive differences associated with brain injuries (Just & Varma, 2007). This cognitive architecture has mainly been validated in medical applications.

ACT-R (Adaptive Control of Thought Rational): is a cognitive architecture (Anderson et al., 2004) that incorporates various specialized modules, including perception modules, which provide mechanisms for real world interaction, and memory modules, for declarative memory (semantic and episodic memory) and for procedural memory (with a production rules set). Among its most remarkable applications are the *Cognitive Tutors for Mathematics*, used in thousands of schools in EEUU.

CHREST (Chunk Hierarchy and REtrieval STructures): is a cognitive architecture that stores its knowledge in a network of *chunks*, interconnected nodes. When the agent learns something new, a new connection of those nodes is created (Gobet, Richman, Staszewski, & Simon, 1997). This cognitive architecture has been validated in chess learning simulations and used in kids' vocabulary improvement scenarios.

CLARION (Connectionist Learning with Adaptive Rule Induction ON-line): is a cognitive architecture whose distinctive feature is the total separation between explicit and implicit knowledge and the way they interact for most tasks (Hélie & Sun, 2010). Clarion has been used for the modeling of cognitive processes such as learning and creativity.

LIDA (Learning Intelligent Distribution Agent): is a cognitive architecture that has a wide variety of computational mechanisms chosen for their plausibility from a psychological point of view (Franklin, 2007). It has been used in intelligent task management systems in military environments and in the control of unmanned autonomous vehicles.

RCS (Real-time Control Systems): is a cognitive architecture developed by the National Institute of Standards and Technology and designed to aid in robot control in laboratory environments (Albus & Barbera, 2005). RCS has evolved to a real-time control architecture for general purpose intelligent systems, particularly in automated factory processes, unmanned vehicles, and military environments.

Soar: is a cognitive architecture designed to model the intelligent behavior of an agent (Laird, 2008). Soar adopts the idea that knowledge, planning, reaction to events, and learning can be integrated in a simple and homogeneous architecture. Soar includes a wide set of mechanisms such as semantic memory, episodic memory, reinforcement learning, and spatial and visual information system. Moreover, this architecture includes an emotion management system designed to influence its reinforcement learning mechanism (Marinier & Laird, 2008). Soar has been used in different applications domains such as control of military planes (Tac-Air Soar) and natural language processing (NL-Soar).

Although most cognitive architectures analyzed lack some essential features for building ECAs (e.g., episodic memory and affective mechanisms), they provide a solid environment for adding this type of requirements. In particular, regarding affective mechanisms, only Soar implements an emotion management system in which emotions influence the reinforcement learning process. However, this emotion system is limited in terms that emotion signals are not used by other components of Soar to influence the global agent's behavior. Moreover, according to the previous analysis, ACT-R and Soar seem appropriate architectures to develop ECAs given their maturity, wide range of features (e.g., different memory systems and learning mechanisms) and available free implementations in different programming languages. We provide in Table 1 a summary of the main characteristics (from the perspective of the development of ECAs) of revised cognitive architectures.

2.2. Affective models

A critical issue in the modeling of affective aspects in ECAs is the lack of consensus about the definition of key concepts such as the term *emotion* (Izard, 2010). According to Gratch and Marsella (2004), *emotions* can be defined as the result of the subjective interpretation of a meaningful event for an agent. In addition, (Clore & Gasper, 2000) demonstrated that emotions influence dif-

[1] http://bicasociety.org/cogarch/architectures.htm.

Table 1
Main characteristics of revised cognitive architectures.

	4CAPS	ACT-R	Chrest	Clarion	LIDA	RCS	Soar
Procedural memory?	Yes	Yes	No	Yes	Yes	Yes	Yes
Semantic memory?	No	Yes	Yes	Yes	Yes (via sparse distributed memory)	Yes	Yes
Episodic memory?	No	Yes	No	Yes	Yes (via sparse distributed memory)	no	Yes (snapshots of working memory)
Reinforcement learning	No	Yes	No	Yes	Yes	No	Yes, for operators (SARSA/Q-learning)
Perceptual memory?	No	Yes	Yes	No	Yes	Yes	Yes (Soar's Spatial Visual System, SSVS)
Affective features?	No	No	No	No	No	No	Some (affects only reinforcement learning)
Defines temporal or capacity constraints?	Yes	No	Yes	No	No	No	No
Implementation	Lisp	Lisp, TCL/Tk, Java	Lisp, Java	Java	Extensible framework in Java	C++, VXworks	C with interfaces to almost any language; Java
(Some) applications	Neuroimage measures, medical applications	Cognitive Tutors, Learning	Chess learning, vocabulary on kids	Simulation of cognitive processes	Task management, autonomous vehicles	Autonomous vehicles	Tac-Air Soar, NL-Soar

ferent aspects such as verbal expressions, non-verbal expressions (e.g., facial gestures and body postures), and cognitive functions (e.g., decision making, and information retrieval). In this context, the literature reports diverse affective models that try to explain and model this influence of emotions in cognition. In this section, we review some representative affective models organized in two groups: *emotion models* and *mixed models*. The first group refers to models that explain the generation and classification of emotions whereas the second group refers to models that combine different aspects such as emotions, mood, personality, empathy, and coping.

2.2.1. Emotion models

The emotion models revised in this section attempt to explain the process of emotion generation and classification mainly based in the concept of *appraisal*. Appraisal-based models of emotion propose that agents constantly evaluate perceived events based on a set of *appraisal dimensions* (Lazarus & Folkman, 1984). In this manner, a particular set of values of those variables are mapped to a single emotion or a range of related emotions. This evaluation process also leads to emotional responses. Appraisal models differ greatly in the number of variables to evaluate and in the mappings between *appraisal dimensions* and generated emotions.

The **OCC model** is the most widely used emotion model for the development of ECAs (Ortony, Collins, & Clore, 1988). This model describes a hierarchy of twenty-two emotion types classified in different categories depending on their relation to external event consequences (joy, guilt), agent actions (pride, reproach), and aspects of objects (love, hate). Each of these categories has different related appraisal variables (e.g., *desirability* and *likelihood*). A simplified version of the OCC model that includes only twelve emotions (Ortony, 2002) (i.e., *joy, hope, relief, pride, gratitude, love, distress, fear, disappointment, remorse, anger and hate*) has also been used in the development of intelligent computational agents (Steunebrink, Dastani, & Meyer, 2009).

The **Roseman's model** proposes two appraisal components for the elicitation of emotions (Roseman, 1996): *Motive consistency* (the evaluation of a situation as inconsistent with the goals of the subject tends to elicit a negative reaction) and *Responsibility/Accountability* (who is responsible of that situation: myself, other agent, or by chance). These values of these appraisal components determine what emotion will be generated. For instance, a good deed done by the agent could generate pride. It was later revised and expanded to generate emotions more similar to the emotions typically generated by human beings in the same situation.

The **Scherer's model** uses up to sixteen different *appraisal dimensions* to elicit emotions (Scherer, 2001). The appraisal follows a complex process based on physiological and psychological aspects: the *multi-level sequential checking model*. This sequence is a step-by-step checking of the sixteen *appraisal dimensions* (e.g., Novelty, Pleasure, Relevance, and Urgency), evaluated on different steps and related with different body systems (neuroendocrine system, autonomous nervous system, somatic nervous system).

2.2.2. Mixed models

The models revised in this section combine emotion mechanisms with aspects such as mood, personality, learning, response selection, coping strategies, and empathy. These models have been implemented and validated in several cases studies. Moreover, most of them have proven useful for the development of ECAs.

FLAME is a model that combines the OCC and Roseman appraisal theories and uses fuzzy logic as part of its behavior selection mechanism (El-Nasr, Yen, & Ioerger, 2000). FLAME includes learning methods to improve its perception of the environment, with associations between objects, events and goals. This model incorporates three components: emotional, decision making and learning components. Its operation cycle is as follows. Events from the environment are evaluated by the emotional component, which calculates emotions and intensities. Afterwards, the learning component modifies the calculated values using the agent's previous experience. These emotions are then filtered and mixed to generate the agent's emotional state, and finally, this emotional state influences the decision making component.

CATHEXIS is a model based on a network of nodes connected to each other (Velásquez & Maes, 1997). Each of these nodes is a *proto-specialist* that represents a class of emotions: *anger, fear, sadness, happiness, disgust and surprise*. Each proto-specialist is connected to sensors that perceive inner and outer stimuli. There are four different types of sensors: neural, sensorimotor, motivational and cognitive. Data from these sensors modify the intensity of the emotions of each proto-specialist, combined with information from other nodes. Finally, a response is selected depending on the set of emotions present in the system.

ALMA is a model based on three layers: emotion, mood, and personality (Gebhard, 2005). Emotions represent an affective response in the short-term and tend to decay quickly once its cause is removed. Their intensity depends on the particular emotion elicited (Kipp, Dackweiler, & Gebhard, 2011), the current mood, and personality (Klesen & Gebhard, 2007). In ALMA, the emotion generation process is based on the OCC model. Mood is defined as a durable emotional state that influences actions from an entity, representing the medium-term. Mood is expressed with three independent traits according to the PAD *temperament* model from Mehrabian (1996). Personality is based on the OCEAN model (McCrae & John, 1992) and defined as the set of mental characteristics of an entity that makes itself unique, representing the long-term.

EMA is a model proposed by Marsella and Gratch (2009) whose purpose is to integrate (1) quick, instinctive emotional reactions to events and (2) further actions due to agent deliberation about its own emotions (*coping*). Its operation cycle is as follows. The agent builds and maintains an ordered sequence of perceived events and the relations between them, regarding its beliefs and goals. Each of these events generates multiple *appraisal frames* that have associated six *appraisal dimensions*. An *appraisal frame* is an extension of an *appraisal* with additional information about the environment. Appraisal frames are then mapped to particular emotions, which are aggregated in a current emotional state and mood. Finally, EMA chooses a coping strategy in response to current emotional state.

Soar-Emote is a model that uses eleven of the sixteen *appraisal dimensions* of the Scherer's model (Marinier & Laird, 2007). It was implemented in Soar (version 9.3). In this model, the emotion of the agent is the set of current values of its *appraisal dimensions*. Mood is a more lasting emotional state, calculated as the average of recent emotions. Finally, mood and emotion are combined to generate a *feeling*, which defines the current emotional state of the agent. This emotional state modifies the reinforcement learning mechanism.

FAtiMA is a model designed to build intelligent agents whose behavior is influenced by their emotions and personality (Dias, Mascarenhas, & Paiva, 2014), This model is aimed at creating believable and empathic agents, as they are more user-friendly and generally they are perceived as more human-like. This model implements a layered model of emotions (reactive and deliberative layers) with integrated planning, learning, and coping mechanisms.

As a summary, it seems that OCC is one of the most widely used model of emotions. However, given its complexity, it is usually simplified and only a few of the 24 emotions proposed are considered in most applications (Ortony, 2002). This model may be combined with other aspects such as mood and personality to reflect long-term emotional states in ECAs, as well as with features such as learning, empathy, and coping to build modern, believable, and complex ECAs.

2.3. ECAs

There is a large volume of literature that explores the applications and development process of ECAs. In this section, we discuss some application domains of ECAs and explain how some of the models reviewed in previous sections have been used in the development of such computational agents.

The affective model ALMA (Gebhard, 2005) has been implemented in ECAs to provide them with emotionally processed information useful to improve their conversational abilities. Particularly, the emotional information generated by ALMA was used to modulate the verbal and non-verbal expressions of these ECAs and inform their selection of dialog and linguistic style strategies. (Reithinger et al., 2006) describe how ALMA is incorporated into the VirtualHuman System, a knowledge-based framework aimed at creating 3D interactive applications for multi-user/agent settings. ALMA allows these computational agents to maintain affective conversations by implementing emotional reactions and expressions.

Similarly, EMA (Marsella & Gratch, 2009) has been incorporated into the architectures of ECAs developed to simulate a diversity of scenarios. In the *The Mission Rehearsal Exercise* project, a virtual-based training program intended to teach soldiers how they should act in stressful situations, EMA was used to influence the decision-making of computational agents and thus allow them to achieve more realistic and human-like behaviors.

The FLAME affective model has been used to dynamically modify the facial expressions of interactive agents (El-Nasr, Ioerger, Yen, House, & Parke, 1999). Moreover, PETEEI (*PET with Evolving Emotional Intelligence*) is an interactive emotional pet that implements FLAME to develop different emotional states (El-Nasr et al., 2000). The simulation of this software pet provides predefined user and pet actions aimed at producing different emotions. The experiments demonstrate that the learning component and the fuzzy logic technique improve the believability of the pet.

The affective model Cathexis (Velásquez, 1997) was included in the architecture of Yuppy, a physical emotional pet capable of displaying certain emotional behaviors according to the particular situation in which it is involved. Yuppy is able to approach people, avoid obstacles, and express emotions. This computational agent has been situated in various controlled environments, demonstrating that Cathexis is an appropriate model for the development of ECAs whose expressions and behaviors are believable.

3. Cognitive-affective architecture

Fig. 1 shows the proposed cognitive-affective architecture. This cognitive-affective architecture is based on Soar (one of the most complete and developed cognitive architectures) and is extended with an affective model inspired by ALMA, combining short-term, medium-term, and long-term affective characteristics (i.e., emotions, mood, and personality, respectively). The three-layer model of emotions defined by ALMA is relatively simple, yet powerful enough and easily integrable in a cognitive architecture like Soar to use it in the development of ECAs. In the next sections, we explain the main modules of the proposed architecture, design and influence of the affective model and finally, the operating cycle of the cognitive-affective architecture.

3.1. Architecture description

The proposed cognitive-affective architecture is composed of several components whose interaction provides ECAs with abilities such as perception, different memory systems, and learning.

The **procedural memory** component stores part of the system knowledge in the form of *production rules* that define the actions that can be performed by the agent. These rules are ordered by a preference system to choose from the available possibilities when several rules are applicable to a given situation. This component also stores elaboration rules that keep updated variables of working memory or calculate information derived from them. This memory system is modified by two mechanisms: (1) *chunking*, which adds new production rules as the result of logical deductions of the system and (2) *reinforcement learning*, which changes the

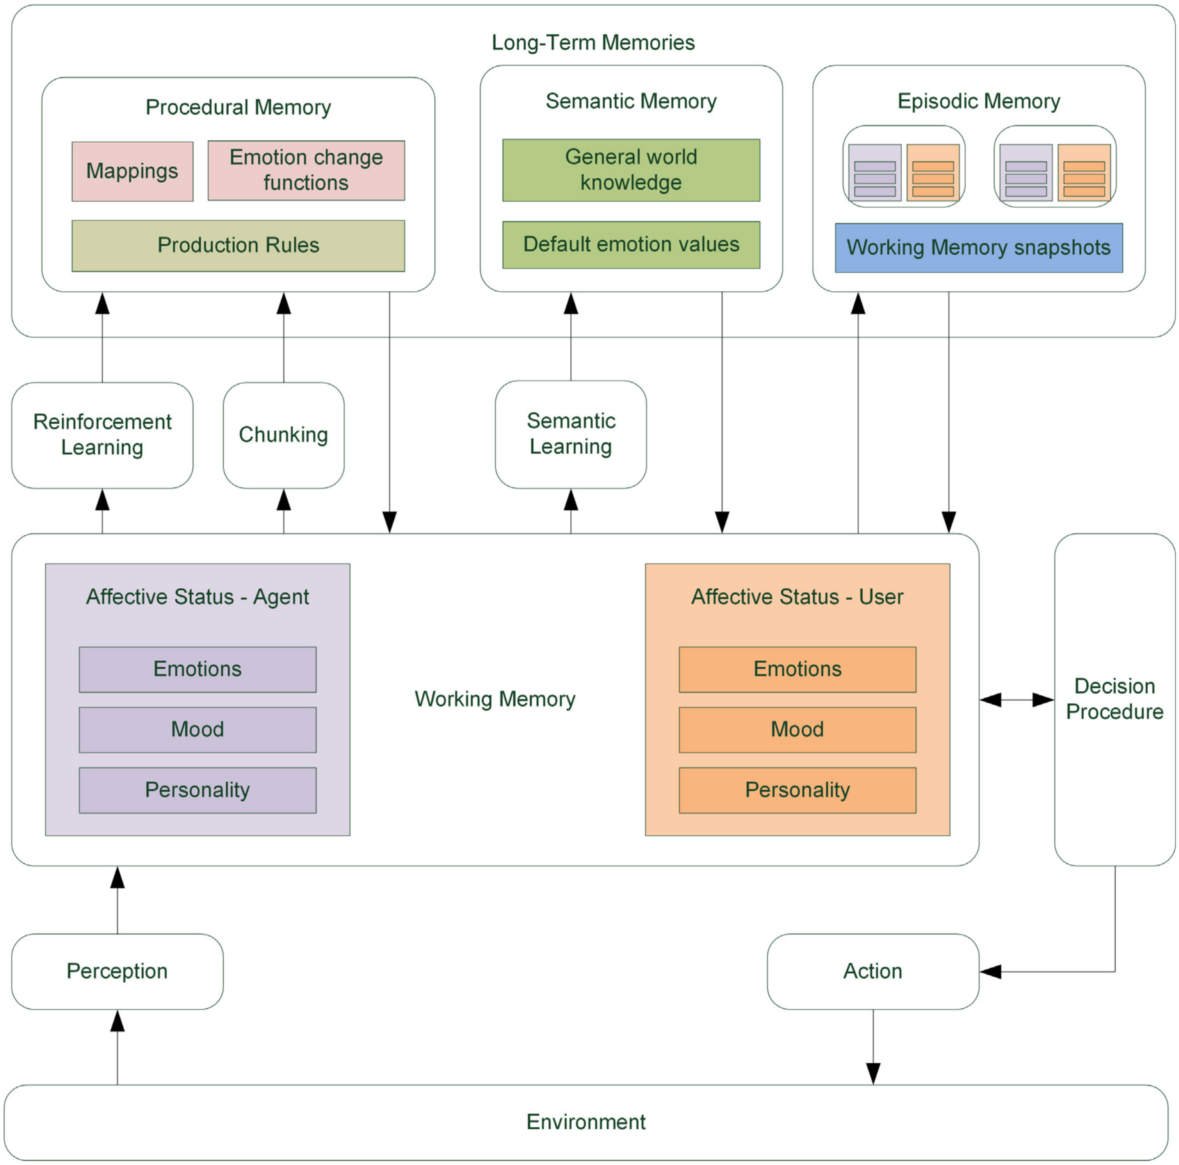

Fig. 1. Proposed cognitive-affective architecture.

preferences associated with the production rules depending on the previous experience of the agent.

The **working memory** component stores information about the current state of the system. It represents a short-term memory that changes according to new information acquired through the input of the system or by agent reasoning mechanisms.

The **semantic memory** component stores information about the world. It is a specialized learning system that stores and maintains up-to-date relevant knowledge.

Furthermore, the **episodic memory** stores memories of past experiences. The information is periodically stored as *snapshots* of working memory in its entirety, sorted by date and time.

The **decision procedure** module decides an action when different options are available. It takes into account information from the other components (e.g., memory systems). It is based on numerical and/or logical preferences.

Finally, the **perception** module constitutes the input interface of the system (supports any type of input) whereas the **action** module performs the action chosen by the decision procedure, possibly changing the environment in some way.

3.2. Affective model

As mentioned above, the proposed cognitive-affective architecture is also based on the affective mechanisms proposed by ALMA (Gebhard, 2005). In this section, we first describe key aspects of these mechanisms and then explain how they are implemented in our proposed model.

In particular, the proposed emotion engine is based on the mechanisms of the emotion model proposed in ALMA, which that takes into account three layers: emotions, mood and personality (which represent short-, medium-, and long-term affective characteristics). The functionality of this proposed emotion engine can be summarized as follows:

- The list of active emotions represents the set of emotions present in the agent, initially empty. The values of the initial mood state (default mood) coincide with the values of the agent's personality.
- Every computed emotion (e.g., joy, hate, and pride) is added to the list of active emotions, encoded as a triad of PAD values,

according to a mapping between emotions and PAD values. Its intensity level depends on the particular emotion, current mood, and personality (Klesen & Gebhard, 2007).
- Emotions decay with time. After a certain period, they are removed from the list of active emotions. This period is emotion-dependent, as some emotions tend to last longer than others.
- The *virtual emotion center* (VEC) is defined as the weighted average of all the emotions present in the list of active emotions, using the Mehrabian's model. The VEC is a point in a 3-D PAD space and its intensity is the module of the vector from the origin $(0,0,0)$ to this point.
- VEC values tend to change erratically, as new emotions appear and the oldest are removed. Therefore, VEC does not represent the current mood, but it is useful to make gradual changes in it, as it changes current mood in a smooth way. VEC *attracts* current mood, with a speed directly proportional to the intensity of VEC.
- The current mood has a tendency to slowly return to the default mood (representing the personality) in the absence of emotions in the system. Also, as personality has PAD values, it's added as a *special emotion* to the active emotions list. Personality does not decay. This addition makes personality more relevant in the whole emotion engine, as it contributes directly to VEC computations and influences the rest of emotions in the system.

Note that the intensity of the VEC (its distance from the origin) is not the same as the intensity of a particular emotion (a value that decays with time according to a mathematical function).

Now we explain how this affective mechanisms are implemented in Soar to support in ECAs features such as empathy and coping as well as to enhance the user experience. We store an *affective state* in the working memory for the agent and another one for each agent or user in the environment. For simplicity, Fig. 1 shows a system with an agent and a unique user. These affective states are an instance of the three-layer emotion model defined by ALMA. In particular, an affective state stores the following elements:

- List of active emotions and their current intensity.
- Virtual emotion center.
- Current mood.
- Personality (default mood).

The stored affective state of other users or agents is an approximate interpretation based on the information that the agent can collect and calculate, and might not match reality.

The following list describes the emotional aspects incorporated to each component of the architecture:

- The semantic memory includes the default numerical values for emotions, as defined by the ALMA model.
- The episodic memory stores information that can be used by the decision procedure to remember past experiences and emotions associated with them.
- The decision procedure can take into account the current emotional state of the system as defined by the affective states or previous agent experience.
- The procedural memory stores *elaboration rules* (red) with mappings between OCEAN and PAD values, and functions for the emotion dynamics that update the affective state of the user and the agent. Additionally, its production rules can take advantage of the new emotional data stored in the architecture.
- The perception module can capture affective information directly from external modules ("*I am sad*") and combine it with internally calculated affective information to enrich the architecture (Ballano Pablo, Baldassarri Santa Lucia, & Cerezo Bagdasari, 2011).

The incorporation of an affective model allows the architecture to manage the behavior of the agent integrating information from different sources (short term memory, episodic memory, semantic memory, preference system) and taking into consideration the affective component in the decision procedure. An agent could react emotionally to the same situation in different ways depending on its previous actions (and consequences) and its current affective state. For example, given the perception of a good action by the user, the agent will probably feel gratitude. The intensity of that gratitude emotion depends on factors such as the specific event, the personality of the agent, its current mood, and any additional information that can be obtained. This approach is similar to the one used in VirtualAgent by Gebhard, Klesen, and Rist (2004) and the biggest difference would lie in the use of the additional possibilities provided by Soar to modify the *appraisal* of events, through its architectural mechanisms. For example, it is possible to change the intensity of an emotion based on the previous experience of the agent to deal with that emotion, and we can use episodic memory to retrieve that information.

3.3. Processing cycle and emotional engine of the architecture

In order to describe the operation of the proposed cognitive-affective architecture, a simple processing cycle of the architecture is explained based on the Soar processing cycle (Laird, 2012, as described in Fig. 2). This processing cycle includes five phases:

- Input phase: In this phase new information is captured by the architecture through its input mechanisms and the values of the components of the affective states are updated.
- Proposal phase: In this phase operators (actions that can be performed) are proposed depending on the state of the system. All possible operators are proposed simultaneously and the architecture is in charge of deciding which one to apply, taking into consideration information from different sources (short term memory, episodic memory, semantic memory, emotional module). Only one operator can be applied per cycle.
- Decision phase: In this phase, one of the operators is chosen to be applied according to the preferences of the system. These preferences take into account other information of the system, such as affective states of the agent and the user, when deciding which operator to apply. It is also stored at this time the reward of the reinforcement learning system as a result of the application of the prior operator.
- Application phase: In this phase the chosen operator is applied. The actions of this operator can update the working memory, retrieve information in other memories and/or provide information to the environment.
- Output phase: In this phase information and orders determined during the application of the operator are sent to the environment.

Running in parallel to the ordinary processing cycle of the architecture shown above, the emotional engine updates periodically the values of the different emotional components of the architecture, at a fixed rate (for instance, 5 s). These operations are performed:

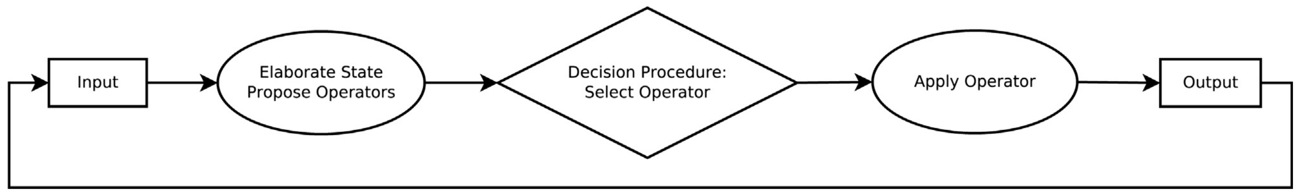

Fig. 2. Soar processing cycle.

```
updateEmotions(){
  updateVirtualEmotionCenterAgent();
  updateVirtualEmotionCenterUser();
  updateCurrentMoods();
}
```

The virtual emotion centers update computes the decay of emotions, and performs a weighted average for the current emotions of the system. This updates the virtual emotion center (PAD values) for the agent and the user. With the new VEC, the current moods for the agent and user can be calculated. We currently use a modified push & pull function as defined by ALMA (Gebhard, 2005). The current mood value *follows* VEC, more quickly the further it is (reflecting more accurately quick emotional changes). Current mood value can never reach VEC, at least theoretically. In practice, VEC values change so quickly that current mood is constantly following VEC, with no inactivity situations.

4. Conclusions

This article presents a study on existing cognitive architectures and affective models, describing their possibilities and making a critical evaluation of their limitations. One of the major problems is the lack of integration between cognitive and affective aspects, essential for up-to-date ECAs. We propose a cognitive-affective architecture based on Soar, extended with an affective model based on ALMA, which allows a holistic approach to the issue. Having all the possibilities of a powerful cognitive architecture being enhanced with affective abilities is a good starting point for the development of state-of-the-art intelligent agents.

This architecture allows to overcome some of the limitations of both Soar and ALMA. Soar is greatly improved by integrating an emotion model. ALMA also benefits from being integrated in a cognitive architecture. Particularly, Soar's episodic memory allows ALMA to have easy access to previous experience of the agent, its emotional state in the past and the experiences related to it, to adjust any kind of characteristics like emotion intensity.

Other approaches have been proposed to address these kind of issues of integrating emotions with a cognitive architecture. For instance, Marinier's work on Soar-Emote (Marinier & Laird, 2008) is certainly remarkable, but emotions only modify the reinforcement learning system, and we believe that the emotional state of the agent should modify its whole behavior to be able to support believable ECAs. Other approach is EMA, built on top of Soar (Marsella & Gratch, 2009). It is a very powerful and comprehensive integration of emotions on a cognitive architecture, mixing together classic AI paradigms (max-utility of states), STRIPS plan representation, BDI model and an appraisal model largely based on OCC variants. However, *power comes with a price*: it is somehow difficult to get the full potential of all these features (Lin, Spraragen, & Zyda, 2012), as it requires the ECA developer or domain expert model to choose some values *ad hoc* (like expected utilities of states) and complicates the integration of EMA with existing ECAs.

The presented architecture is simpler, yet powerful enough to offer a wide range of features. One of its main goals is to be easily integrable with previous systems, and to provide cognitive and affective capabilities to an agent without unwanted interference with its performance.

We are now integrating the proposed architecture in the VOX-System (Serón & Bobed, 2016), a system developed by the GIGA Affective Lab[2] that explores the synergies between the world of the ECAs and semantic information. It is based on the use of ontologies and logical reasoners which use description logic, allowing an ECA to be enhanced with knowledge-related capabilities. We expect that the integration of a cognitive-affective architecture with emotional mechanisms that can influence the behavior of the agent will increase the realism and believability of the ECA, and improve the user experience.

Acknowledgments

This work was supported by a grant from the Spanish "Dirección General de Investigación", contract number TIN2015-67149-C3-1R.

References

Albus, J. S., & Barbera, A. J. (2005). RCS: A cognitive architecture for intelligent multi-agent systems. *Annual Reviews in Control, 29*, 87–99.
Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111*, 1036–1060.
Ballano Pablo, S., Baldassarri Santa Lucia, S., & Cerezo Bagdasari, E. (2011). *Evaluación de un sistema multimodal de reconocimiento de emociones*. Master's thesis.
Cassell, J., & Bickmore, T. (2000). External manifestations of trustworthiness in the interface. *Communications of the ACM, 43*, 50–56. http://dx.doi.org/10.1145/355112.355123.
Clore, G. L., & Gasper, K. (2000). Feeling is believing: Some affective influences on belief. In N. H. Frijda, A. S. R. Manstead, & S. Bem (Eds.), *Emotions and beliefs: How feelings influence thoughts* (pp. 10–44). Cambridge, UK: Cambridge University Press.
Damasio, A. R. (1994). Descartes' error: Emotion, reason, and the human brain. *1st ed.*. New York: Putnam Grosset Books.
Dias, J., Mascarenhas, S., & Paiva, A. (2014). Fatima modular: Towards an agent architecture with a generic appraisal framework. In *Emotion modeling* (pp. 44–56). Springer.
El-Nasr, M. S., Ioerger, T. R., Yen, J., House, D. H., & Parke, F. I. (1999). Emotionally expressive agents. In *CA '99: Proceedings of the computer animation* (pp. 48). Washington, DC, USA: IEEE Computer Society.
El-Nasr, M. S., Yen, J., & Ioerger, T. R. (2000). Flame: Fuzzy logic adaptive model of emotions. *Autonomous Agents and Multi-Agent Systems, 3*, 219–257. http://dx.doi.org/10.1023/A:1010030809960.
Franklin, S. (2007). A foundational architecture for artificial general intelligence. In *Proceedings of the 2007 conference on advances in artificial general intelligence: concepts, architectures and algorithms: proceedings of the AGI workshop 2006* (pp. 36–54). Amsterdam, The Netherlands: IOS Press. <http://dl.acm.org/citation.cfm?id=1565455.1565460>.
Gebhard, P. (2005). A layered model of affect. In *4th International joint conference of autonomous agents & multi-agent systems (AAMAS'05)* (pp. 29–36). ACM Press.
Gebhard, P., Klesen, M., & Rist, T. (2004). Coloring multi-character conversations through the expression of emotions. In *Affective dialogue systems* (pp. 128–141). Berlin/Heidelberg: Springer. <http://www.springerlink.com/content/aq1mp79b30hrgket>.

[2] http://giga.cps.unizar.es/affectivelab/.

Gobet, F., Richman, H., Staszewski, J., & Simon, H. (1997). Goals, representations and strategies in a concept attainment task: the EPAM model. *The Psychology of Learning and Motivation, 37*, 265–290.

Gratch, J., & Marsella, S. C. (2004). A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research, 5*, 269–306. <http://ict.usc.edu/pubs/A%20Domain-independent%20Framework%20for%20Modeling%20Emotion.pdf>.

Hélie, S., & Sun, R. (2010). Incubation, insight, and creative problem solving: a unified theory and a connectionist model. *Psychological Review, 117*, 994–1024. http://dx.doi.org/10.1037/a0019532. <http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=20658861&retmode=ref&cmd=prlinks>.

Izard, C. E. (2010). The many meanings/aspects of emotion: Definitions, functions, activation, and regulation. *Emotion Review, 2*, 363–370. http://dx.doi.org/10.1177/1754073910374661.

Just, M. A., & Varma, S. (2007). The organization of thinking: What functional brain imaging reveals about the neuroarchitecture of complex cognition. *Cognitive, Affective, & Behavioral Neuroscience, 7*, 153–191. <http://www.ingentaconnect.com/content/psocpubs/cabn/2007/00000007/00000003/art00001>.

Kasap, Z., & Magnenat-Thalmann, N. (2007). Intelligent virtual humans with autonomy and personality: State-of-the-art. *Intelligent Decision Technologies, 1*, 3–15. <http://dl.acm.org/citation.cfm?id=2595898.2595900>.

Kipp, M., Dackweiler, T., & Gebhard, P. (2011). Designing emotions. *KI - Knstliche Intelligenz, 25*. http://dx.doi.org/10.1007/s13218-011-0110-2.

Klesen, M., & Gebhard, P. (2007). Affective multimodal control of virtual characters. *The International Journal of Virtual Reality, 6*, 43–54.

Laird, J. E. (2008). Extending the soar cognitive architecture. In P. Wang, B. Goertzel, & S. Franklin (Eds.). *Frontiers in artificial intelligence and applications* (Vol. 171, pp. 224–235). IOS Press. AGI <http://dblp.uni-trier.de/db/conf/agi/agi2008.html#Laird08>.

Laird, J. E. (2012). *The soar cognitive architecture*. The MIT Press.

Lane, R. D., Nadel, L., Allen, J. J. B., & Kaszniak, A. W. (2000). The study of emotion from the perspective of cognitive neuroscience. In R. D. Lane & L. Nadel (Eds.), *Cognitive neuroscience of emotion*. New York: Oxford University Press.

Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research, 10*, 141–160. http://dx.doi.org/10.1016/j.cogsys.2006.07.004.

Lazarus, R., & Folkman, S. (1984). *Stress, appraisal, and coping*. Springer Publishing Company. <https://books.google.nl/books?id=i-ySQQuUpr8C>.

Lin, J., Spraragen, M., & Zyda, M. (2012). Computational models of emotion and cognition. *Advances in Cognitive Systems*, 59–76.

Marinier, R., & Laird, J. E. (2008). Emotion-driven reinforcement learning. *Cognitive Science, 25*, 115–120.

Marinier, R. P., & Laird, J. E. (2007). Computational modeling of mood and feeling from emotion. In *Proceedings of the 29th annual conference of the cognitive science society* (pp. 461–466). Cognitive Science Society.

Marsella, S. C., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research, 10*, 70–90. <http://ict.usc.edu/pubs/EMA-%20A%20process%20model%20of%20appraisal%20dynamics.pdf>.

McCrae, R. R., & John, O. P. (1992). An introduction to the five-factor model and its applications. *Journal of Personality, 60*, 175–215. http://dx.doi.org/10.1111/j.1467-6494.1992.tb00970.x.

Mehrabian, A. (1996). Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in temperament. *Current Psychology, 14*, 261–292. http://dx.doi.org/10.1007/BF02686918.

Ortony, A. (2002). On making believable emotional agents believable. In R. Trappl, P. Petta, & S. Payr (Eds.), *Emotions in humans and artifacts* (pp. 189–211). MIT.

Ortony, A., Collins, A., & Clore, G. L. (1988). *The cognitive structure of emotions/Andrew Ortony, Gerald L. Clore, Allan Collins. (Pbk. ed. ed.)*. Cambridge [England]; New York: Cambridge University Press. <http://www.loc.gov/catdir/toc/cam026/87033757.html>.

Phelps, E. A. (2006). Emotion and cognition: Insights from studies of the human amygdala. *Annual Review of Psychology, 57*, 27–53.

Picard, R. W. (1997). *Affective computing*. Cambridge, MA, USA: MIT Press.

Reisenzein, R., Hudlicka, E., Dastani, M., Gratch, J., Hindriks, K., Lorini, E., & Meyer, J.-J. C. (2013). Computational modeling of emotion: Toward improving the inter- and intradisciplinary exchange. *IEEE Transactions on Affective Computing, 4*, 246–266.

Reithinger, N., Gebhard, P., Löckelt, M., Ndiaye, A., Pfleger, N., & Klesen, M. (2006). Virtualhuman: Dialogic and affective interaction with virtual characters. In *ICMI '06: Proceedings of the 8th international conference on multimodal interfaces. New York, NY, USA* (pp. 51–58).

Rodríguez, L.-F., & Ramos, F. (2015). Computational models of emotions for autonomous agents: Major challenges. *Artificial Intelligence Review, 43*, 437–465.

Roseman, I. J. (1996). Appraisal determinants of emotions: Constructing a more accurate and comprehensive theory. *Cognition & Emotion, 10*, 241–278.

Scherer, K. R. (2001). Appraisal considered as a process of multi-level sequential checking. In A. Schorr, K. R. Scherer, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 92–120). New York and Oxford: Oxford University Press.

Scherer, K. R., Bänziger, T., & Roesch, E. (2010). *A blueprint for affective computing: A sourcebook and manual*. Oxford University Press.

Scheutz, M. (2004). Useful roles of emotions in artificial agents: a case study from artificial life. In *AAAI'04: Proceedings of the 19th national conference on artificial intelligence* (pp. 42–47). AAAI Press.

Serón, F. J., & Bobed, C. (2016). Vox system: A semantic embodied conversational agent exploiting linked data. *Multimedia Tools and Applications, 75*, 381–404.

Steunebrink, B. R., Dastani, M., & Meyer, J.-J. C. (2009). The OCC model revisited. In *Proceedings of the 4th workshop on emotion and computing*. .

Thibadeau, R., Just, M. A., & Carpenter, P. A. (1982). A model of the time course and content of reading. *Cognitive Science, 6*, 157–203. http://dx.doi.org/10.1207/s15516709cog0602_2.

Velásquez, J. D. (1997). Modeling emotions and other motivations in synthetic agents. In *Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence* (pp. 10–15). Providence, Rhode Island: AAAI Press.

Velásquez, J. D., & Maes, P. (1997). Cathexis: A computational model of emotions. In *Proceedings of the first international conference on autonomous agents AGENTS '97* (pp. 518–519). New York, NY, USA: ACM. http://dx.doi.org/10.1145/267658.267808.