The TRIPLE Cognitive Architecture: Implementation of Embodied Agents Based on a Cognitive Model

Maurice Grinberg (mgrinberg@nbu.bg), Vladimir Haltakov (vhaltakov@gmail.com),
Hristo Laftchiev (laftchiev@gmail.com), Stefan Kostadinov (stefan@yobul.com)
Central and Eastern European Center for Cognitive Science, New Bulgarian University,
Montevideo 21, 1618 Sofia Bulgaria

Abstract

This paper introduces a new model which is intended to combine the flexibility of a cognitive architecture and the power new technologies developed for using semantic knowledge in the WEB. The cognitive architecture (called TRIPLE) is hybrid integrate three modules which run in parallel: a reasoning, a connectionist and an emotion engines. The reasoning engine deal with the current task, processes the perceptual input to the system, and interacts with the memory of the architecture and the other two engines. The reasoning engine can augment by inference and consistency checks the representation of the task and the content of the working memory. The latter is determined by a spreading of activation mechanism performed by the connectionist engine which is also responsible for retrieval of knowledge from memory and mapping it to the task. This paper focuses on this latter module and presents its principles and implementation.

Keywords: Cognitive architectures, embodied conversational agents.

Introduction

Virtual environments like the Internet become more and more complex and rich and comparable with real environment with respect to the amount of information and its complexity. As is the case in robotics, human level of performance is difficult to achieve especially with respect to the level of communication, perception, etc. At the same time, although far from human-like performance the cognitive architectures can come supposedly closer to the way humans behave in complex environments (see e.g. Sun, 2006) in comparison with traditional AI approaches. Thus they seem to be good candidates to be used as 'minds' for Embodied Conversational Agents (ECA) living in virtual environments (Krenn, 2008). Moreover cognitive architectures are expected to outperform AI counterparts (e.g. BDI architectures) by better coping with problems related to perception and action on one hand and the interaction with human users on the other. Both require context sensitivity, flexibility, believability, and personalization of the ECA behavior. However, acceptable performance in real world settings poses the question of scalability of the cognitive architectures. The latter is an important constraint and a powerful motivation for model improvements in efficiency but preserving cognitive plausibility and advantages at the same time.

The work presented here has been inspired and determined by the above mentioned consideration. The cognitive architecture TRIPLE, presented here, has been the result of attempting to adapt a research cognitive architecture for a multi-agent platform. The result is a new architecture which is based on modern semantic web technologies but implements cognitively plausible mechanisms. In our opinion, the requirements of efficiency in terms of time and representational power can be viewed as belonging to the pluralistic approach to cognitive phenomena (see e.g. Dale, 2008).

The model TRIPLE, introduced for the first time in (Grinberg & Kostadionv, 2008), has been designed to be used in artificial cognitive model and for ECA platforms in particular. On one hand, TRIPLE includes all the necessary mechanisms to be a fully fledged cognitive architecture and on the other its implementation has been subject to maximal computational efficiency requirements in order to allow real time functioning. These two constraints lie at the basis of the architecture: adding all the useful cognitive modeling techniques which allow flexibility, context sensitivity and individuality of the agent and at the same time achieving maximal computational optimization of the implementation and including for instance very efficient inference methods (Kiryakov, Ognyanoff, & Manov, 2005).

Following this strategy, the model has been designed in three parts that function in parallel. The so-called Reasoning Engine (RE) is coordinating and synchronizing the activities of the model and relates the agent with its environment and with the tools it can use in this environment like communicate with a user, access and retrieve knowledge from ontologies, etc. (see Grinberg & Kostadionv, 2008). RE is also responsible for instance learning – storing of useful episodes in LTM after evaluation. Part of RE is the Inference Engine (IE) which operates on a limited amounts of active relevant knowledge (the most active part of the WM of the agent). Its main role is to augment parts of WM with inferred knowledge and do consistency checks.

The second module of Triple is the so-called Similarity Assessment Engine (SAE). It is designed to be a connectionist engine, based on fast matrix operations and is supposed to run all the time as an independent parallel process. The main mechanism is activation spreading in combination with similarity or correspondence assessment mechanisms which allow retrieval of knowledge relevant to the task at hand. The communication of SAE with RE is based on events related to the level of confidence for a match between the task and WM content. The information retrieved can correspond to the input and goals of the architecture at different level of abstraction base on flexible similarity assessment mechanisms thus allowing for case based reasoning and analogy. The main principle behind SAE (and the main novelty in TRIPLE) is the formation of various distributed representations based on symbolic representation of knowledge. Such distributed representations encode different aspects of the symbolic knowledge and allow for connectionist type of processing with all its advantages. For instance, such complementary representations can be built by using the taxonomy relations in the semantic memory and co-occurrence in episodic memory, of latency semantic analysis (see Ramscar & Yarlett (2003), for a related approach).

The third important part of the architecture is the Emotion Engine (EE) which is based on the FAtiMA emotional agent architecture (Dias & Paiva, 2005). FAtiMA generates emotions from a subjective appraisal of events and is based on the OCC cognitive theory of emotions (Ortony, Clore, & Collins, 1988). EE, similarly to SAE, is supposed to run in parallel and influence various parameters of the model like the volume of WM, the speed of processing, etc. (see Vankov, Kiryazov, & Grinberg (2008) for a simple exploration of the role of emotions on analogy making). The availability of the emotional engine allows for higher believability and usability based on the emotional expressions and gestures corresponding to the current emotional state of the agent. The importance of emotions for ECAs and the interaction of users with them have been pointed out for instance in Becker, Kopp, & Wachsmuth (2007).

The TRIPLE model is connected to the DUAL/AMBR architecture (Kokinov, 1994) by inheriting some important mechanisms. TRIPLE, similarly to DUAL/AMBR, makes use of spreading of activation as a method for retrieval from memory of the most relevant episodic or general knowledge. The mapping of the knowledge retrieved to the task at hand and to the current input to the system is based on similarity and analogy in both models (Kiryazov et. al., 2007; Kostadinov, Petkov, & Grinberg, 2008; Kostadinov & Grinberg, 2008). However the underlying mechanisms are essentially different. In DUAL/AMBR knowledge representation is based on a large number of micro-agents which perform local, decentralized operations and are dualistic in the sense that they spread activation and perform symbolic operations at the same time. The messages exchanged between the micro-agents (marker passing) trigger the establishment of mappings and structural correspondence assessment. The speed of the symbolic processing of the messages depends on their level of activity over time. In TRIPLE, an attempt has been made to achieve the same and better functionality on the basis of clearly separated symbolic mechanisms (reasoning, inference, consistency checks, anticipation, etc.) and connectionist mechanisms (spreading of activation over different types of connections, similarity assessment based on distributed representations and activation, etc.). A third component is the emotional module (EE) which is missing in DUAL/AMBR(Kiryazov & Grinberg, 2009).

In DUAL/AMBR the 'duality' is achieved at the level of each micro-agent while in TRIPLE it is achieved by two systems which run in parallel and communicate on an event-driven basis. An important additional difference is that TRIPLE is using a fully fledged reasoning part in the standard AI sense, which is not available in DUAL/AMBR. The inference and entailment capabilities are integrated with the spreading of activation and evaluation of retrieval and action planning. Only the most active part of WM, corresponding to the focus of attention of the system is subject to augmentation based on inference and to other symbolic processing like evaluation, transfer, and action. The Amine platform (Kabbaj, 2006) has similar augmentation mechanisms which are based on purely symbolic manipulation and are not conditioned by any attention mechanism of the system (see the discussion of 'elicitation' and 'elaboration' in Kabbaj (2006)).

In this paper we will focus on the presentation of the cognitive architecture TRIPLE by describing the structure and functioning of the three main engines, mentioned above and their interaction. The main focus will be on the connectionist part (SAE) which we considered as the main contribution of the new architecture.

The TRIPLE Cognitive Architecture

As stated earlier, TRIPPLE has been designed to be applied as underlying cognitive architecture for an embodied cognitive agent (a robot or a virtual agent). The general conceptualization of the interaction between the cognitive architecture (‘Mind’) and the ways available to the agent of perceiving and acting (‘Body’) in the environment, is shown in Figure 1. The advantage of this representation is the possibility to consider the cognitive architecture part of both physically and virtually embodied agents on equal footing.

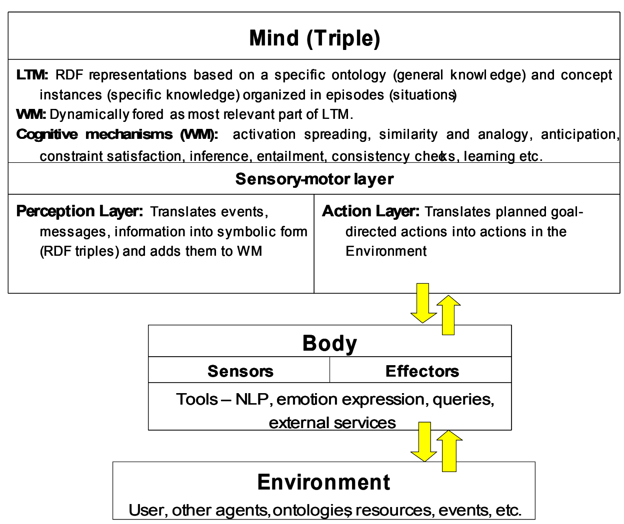

Figure 1: A Mind-Body-Environment conceptualization of the TRIPLE cognitive architecture as ‘Mind’ in a embodied cognitive agent platform.

Figure 1 shows the ‘Mind’ with its specific knowledge structures and mechanisms, the ‘Sensory-Motor Layer’, making a mediated connection between the ‘Mind’, and the ‘Body’ (consisting of various ‘Tools’, including ‘sensors’ and ‘effectors’). The sensors and effectors connect the ‘Sensory-Motor’ layer to the ‘Environment’ and implement perception and action in it.

The ‘Sensory-Motor’ layer provides symbolically represented knowledge to the ‘Mind’ and thus makes a mediated connection between the ‘Mind’ and the ‘Environment’. This is needed as far TRIPLE is a hybrid cognitive architecture and knowledge is represented in symbolic form.

A schematic presentation of the TRIPLE cognitive architecture is given in Figure 2. And the way it operates is the following. All the knowledge is organized as a semantic network in which there is general knowledge expressed as concepts (including relations) and instances of concepts. The latter are used to store episodes which form the episodic memory. All knowledge to become accessible to the architecture must be expressed as sub-classes and/or instances of existing concepts in LTM (i.e. in symbolic form). The communication with the environment is organized as input to the system (via the sensors of the ‘Body’) in which the task and the goal are described and attached to the LTM by RE. All concepts and instances of concepts are considered to be nodes of a neural network with the task nodes as sources of activation as in the DUAL/AMBR architecture (Kokinov, 1994). Whenever there is input to the system activation is spread throughout LTM and relevant nodes start to become activated. The level of activation is considered to be a measure of the relevance of the activated knowledge to the input presented. The most active part in LTM is considered to be the WM of the system. The SAE is running all the time and in parallel trying to establish correspondences between the input and the output thus looking for past episodes or knowledge that could help the completion of the task. The correspondence can be established based on similarity, analogy, semantic relation, or on any other meaningful principle. The higher the level of correspondence is, the higher the probability of attracting the attention of the system is, in which case RE retrieves the correspondence hypothesis and starts the evaluation process. The processing in RE is limited to the most active part of WM is serial and is considered to correspond to processing which require attention and awareness. Such are the inference and consistency check processes which are part of the evaluation and transfer made by RE. If the evaluation is successful additional knowledge is transferred from long term memory which is evaluated on its turn by RE. If a correspondence contradicts previous knowledge or is impossible it is suppressed. If for some reasons no appropriate knowledge in LTM is retrieved RE starts a process of augmentation of the task description by using inference and entailment. This allows for more flexibility in the retrieval and encoding independence. Once a good candidate for task solution has been retrieved and evaluated it is transformed into appropriate actions and sent to the ‘Body’ for execution.

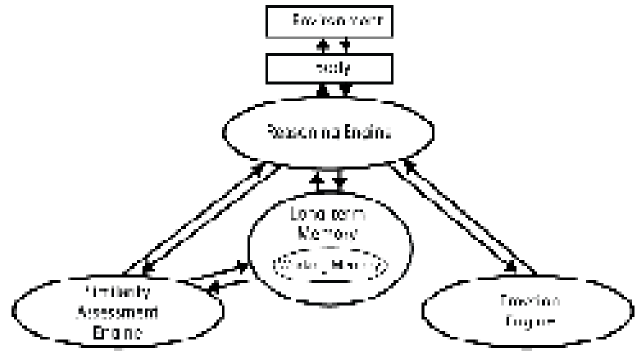

Figure 2: Structure of the TRIPLE cognitive architecture.

The Similarity Assessment Engine (SAE)

General Approach As stated above SAE is the most distinguishing feature of TRIPLE. It is a connectionist engine based on distributed representations of the elements of LTM and of the input to the system. The main principles behind SAE are the following:

- Creation of a network of connections (weight matrices) based on some principle like association, taxonomy, participation in relations and actions, semantic relatedness, etc. and spreading of activation over these networks (in this way there could be different types of activation depending on the building principles of the networks;
- Similarity assessment based on distributed representations obtained from the symbolic representation of knowledge (e.g. the matrix of weights of the taxonomical relations for a concept or instance of concept, or the weights to relations and actions in which they participate).

The idea behind these two principles can be related to a kind of reverse engineering approach. It starts from symbolic encoding of knowledge in a semantic network, in episodes, or even as implicitly present in texts from a domain. Then the task is to build distributed representations out of the structured symbolic representation. For instance, a taxonomy of the type ‘a kitchen chair’ → ‘sub-class-of’ → ‘chair’ → ‘sub-class-of’ → ‘furniture’ → etc., can be seen as a distributed representation of ‘kitchen-chair’ over the members of the hierarchy ‘chair’, ‘furniture’, etc., and represented by the giving weights to the connection of ‘kitchen-chair’ with any of the higher lying concepts (e.g. by a distance function). Similarly, ‘chair’ can be represented in terms of the relations and action in which it participates like ‘being ON the floor’, ‘being seated ON’, ‘being BEHIND a table’, etc. which will give a distributed representation of objects and instances of objects over relations and actions. As far as the whole knowledge of the cognitive architecture is represented in the semantic net (or ontology) any distributed representation of knowledge elements, obtained by the method outlined above, will be in terms of elements of the same semantic network. In this way, the knowledge of the system can be represented as a multi-dimensional space, spanned by all elements in LTM (plus the instances in episodic memory) and using the various types of connections among the elements various representations can be built of sub-spaces of the full knowledge space in terms of other or the same sub-spaces. This is consistent with a transformational perspective about the relations among objects. However, this static picture is made dynamic by the activation spreading mechanism. Effectively only the active elements in LTM, which characterize the state of the cognitive system at a specific moment in time, should contribute to the distributed representation. Thus, this approach to building distributed representation is highly dynamic and context sensitive. For instance, 'kitchen-chair' and 'kitchen-table' can be highly dissimilar if only the lowest nodes in the taxonomy are activated, e.g. 'chair' and 'table', respectively. But when 'furniture' gets active they will have something in common and will become more similar as activation spreads higher in the taxonomy. Thus in assessing similarity or more generally correspondence only the active nodes matter and the correspondence measure varies with time reflecting the context of the situation as given by the activation patterns.

The construction of such distributed representation is related to the task considered. In a case-based reasoning framework, taxonomic connections can be enough to construct useful distributed representation. For analogy reasoning – the relational structure has to be also taken into account.

Implementation

The implementation of SAE in TRIPLE follows the principles laid down above. In the present implementation there is only one activation spreading mechanism (for activation spreading based on semantic similarity see Ramscar & Yarlett, 2003) and two distributed representations. The first is based on taxonomic relations of the type 'instance-of' and 'sub-class-of'. The second takes into account the participation in realtions (including actions). The aim of the present implementation is to investigate what are the capabilities of TRIPLE in case-based type of reasoning and analogy.

As described above retrieval from memory is based on the level of correspondence of the task elements to knowledge in LTM. The measure of correspondence chosen is similarity of task and LTM elements measured on the basis of the two distributed representations by normalized scalar product. The whole process involves three subprocesses running in parallel – activation spreading, correspondence assessment and constraint satisfaction.

Activation spreading takes place by using the equation:

$$a^t = f_a(W_{LTM+T}a^{t-1}) \tag{1}$$

where *a is* a vector, corresponding to the activation, represented in the space of all nodes of LTM and the task nodes; f_a is a standard activation spreading function whicih keeps activity in the range [0,1]; W_{LTM+T} is aweight matrix with elements corresponding to weights of links between the nodes in LTM (chosen equal to 1); t is the current iteration.

The distributed representations of any memory or task element are presented in matrices, called similarity matrices (denoted X). The elements to be compared are represented as rows in X and the normalized scalar product of rows gives a correspondence or similarity measure for any two elements, weighted by the current activation:

$$S_{ij}^t = a_i^t a_j^t \sum_k X_{ik} X_{jk} a_k^t, \tag{2}$$

where $a^t{}_i$ and $a^t{}_j$ are the activities of the elements compared at the time t; $a^t{}_k$ are the activities of the elements which form the 'basis' of the distributed representation of elements i and j. The matrix elements X_{ik} can be the weights of the connections between the corresponding elements or be evaluated using a distance functions (see the discussion bellow). The inclusion of activity in the similarity evaluation makes eq. (2) dynamic and dependant only on the most active elements in LTM and in the input.

Once the similarities have been computed they become external input in a Constraint Satisfaction Network (CSN) based on the IAC model of McClleland & Rumelhart, (1981). A similar approach has been explored in Ming Mao et al. (2008). In this sub-network the similarity is considered to be a special kind of activity standing for probability of mapping between two elements characterizing a correspondence hypothesis. Interaction and competition take place among such hypotheses for correspondence based on the constraint to have only one mapping between a task element and a memory element and the requirement to have consistent mappings between connected elements (elements which are connected themselves to correspond to connected elements). In this way, each correspondence hypothesis supports any other non-competing correspondence hypothesis. In other words the hypotheses for correspondence of connected elements will support each other if they do not map both to the same element (e.g. 'chair' → 'table' and 'on'→'in-front', coming from episodes in which 'the chair is on the floor' and 'the table is in-front of the window'). The latter is implemented by a connectivity matrix for LTM which contains weights inversely proportional to the distance between two elements, excluding taxonomy relations like 'instance-of' and 'sub-class'.

The implementation of this mechanism is done as follows. A matrix H is initialized with dimensions equal to the similarity matrix between task and LTM elements. Then it is updated by after each cycle of activation spreading and similarity calculation by the formula:

$$H_{ij}^t = dH_{ij}^{t-1} - c_{inhib}H_{ij,inhib}^{t-1} + c_{excit}H_{ij,excit}^{t-1} + (1-d)S_{ij}^t, \tag{3}$$

where i refers to a task element, and j to a LTM element. Every cell (i,j) of the matrix H represents a hypothesis between the elements i and j, with probability of being true estimated by the value of H_{ij}. The coefficients in Eq. (3) are chosen in such a way so that H_{ij} remain in the interval [0,1]. The terms $H_{ij,\ inhib}$ and $H_{ij,\ excit}$ are the inhibitory and the excitatory contributions to H_{ij}, respectively. The quantity d is a decay parameter. The parameters c_{inhib} and c_{excit} are chosen to so that the values of H_{ij} are kept within the interval [0,1].

The quantities $H_{ij,\ inhib}$ and $H_{ij,\ excit}$ are calculated by using the following expressions:

$$H_{ij,inhib} = \sum_{t \neq i} H_{it} + \sum_{m \neq j} H_{mj} \qquad (4)$$

$$H_{ij,excit} = \sum_{\substack{t,m \\ t \neq i \\ m \neq j}} H_{tm} (C_{jm} + C_{it}) \qquad (5)$$

where H_{it} and H_{mj} are probabilities of correspondence from the same row and column in the matrix H as the element H_{ij} and thus compete with it because i-th row contains other correspondence hypotheses with memory element i, and the j-th column other correspondence hypotheses involving memory element j. In, eq. (5), the summation goes over all the hypotheses outside the row i and the column j multiplied by the connectivity weights discussed above. The quantity C_{jm} and C_{it} account for the closeness in terms of connections of any memory element m to the memory element j, and of task element t to task element i. The connectivity weights are calculate using the function C_{pk}=1/n, where n is the number of connections to be followed in order to reach p from k. (e.g. 1 if there is one direct connection between them, 1/2 for two connections, etc.)

The Reasoning Engine (RE)

The reasoning engine integrates the results of the operation of the SAE and EE on one hand and the communication flow with the 'Body' and the 'Environment' on the other (see Figure 2). It works with the sensory-motor part of TRIPLE by receiving the information from the sensors (user utterances, information about results from actions, etc.) and sending action commands to the 'Environment'. The main interactions with the 'Body' and 'Environment' are the same as the ones reported in Kostadinov & Grinberg (2008) and are only given below for completeness.

The task set of elements (instances of concepts and relations describing the task and the goal) is the source of activation for the SAE. As explained above SAE runs in parallel (together with the EE) and has continuously information about possible candidates for retrieval and consideration by RE. The latter evaluates the level of similarity and based on that establishes candidate correspondences between the elements of the task and LTM.

When the correspondences between the task and LTM are established the RE verifies and evaluates them and eventually rejects or confirms them. Based on the existing correspondences, parts of past episodes are transferred as possible candidates for task completion. These transferred memory elements are evaluated on their turn for consistency with task until eventually an action transfer is chosen and the appropriate action is executed. We assume that RE must work serially because it deals with high level reasoning tasks. It actually deals with the most active part of WM which can be considered to be the focus of attention of the system.

If the task is considered completed (e. g. a question is answered) the whole episode with the task and its completion is stored in LTM as an experience episode for future use. Any new knowledge acquired during a task completion base on inference or knowledge retrieval from an external source is stored as general knowledge.

In order to achieve the above mentioned tasks, RE incorporates the following processors:

- Similarity Assessment Processor: handles the initial similarity assessments from SAE, eliminates the inconsistency and establishes correspondence hypotheses;
- Correspondence Processor: processes the correspondence hypotheses. The most probable correspondence hypotheses (as determined by the CSN mechanisms described above) are used as a baseis for retrieval and transfer;
- Transfer Processor: removes inconsistent transferred pieces of knowledge and evaluates transferred knowledge with respect to its consistency with the current goals. This is the module that actually adds transferred knowledge to WM. Part of it contains possible actions to perform.
- Action Processor: receives a list of possible actions to execute from the Transfer Processor. Contradictory actions are evaluated, based on their activation and the one with the highest activation is executed.

The Emotional Engine (EE)

The Emotional Engine works in parallel with the SAE and provides continuous fine-tuning of the rest of the cognitive platform (Kiryazov & Grinberg, 2009). The so-called fine-tuning is done by the constantly changing emotional state. The effects of EE on the processing in TRIPLE, are related for instance to volume of WM, the speed of the inference mechanisms, the interaction between RE and SAE, and SAE itself, the severity of the reasoning constraints, etc. Another major influence is the tagging with emotional tags of episodes in LTM. This allows for retrieval of mood congruent episodes. A detailed description of EE will be available in Dias et al. (2009). Some exploration in this direction, based on simulations have been reported in Vankov, Kiryazov, & Grinberg (2008).

Discussion and Conclusion

In the paper, a new cognitive architecture, called TRIPLE was described. The architecture was developed and is used in a multi-agent Web based platform. However the main focus of this paper was to present the main components of the architecture and to focus on a connectionist module of the architecture (SAE) which is based on spreading of activation and distributed representations derived from symbolic localist representation like semantic networks.

The main idea behind the architecture is to combine a cognitive model which would bring flexibility, context sensitivity based on activation spreading and episodic memory highly efficient AI approaches like logical inference over ontologies. At the core of the model is the combination and simultaneous action of a symbolic RDF triple based inference engine and a fast linear algebra-based connectionist engine. The main principle behind the connectionist engine is to combine activation spreading and similarity assessment based on dynamic distributed representations in order to select a tractable amount of LTM knowledge most relevant to the task. The architecture is then able to make inferences on that part, augment it and increase the efficiency of the system by discarding timely contradictory or inconsistent information.

The three main 'engines' of the model – the reasoning, the inference, and the similarity assessment engines – have been already implemented and parts of them tested. Although the results are very promising they are too preliminary to be reported here and will be published in a next paper.

Acknowledgments

We would like to acknowledge the partial support for this work of EC FP6 project RASCALLI.

References

Becker, C., Kopp, S., & Wachsmuth, I. (2007). Why emotions should be integrated into conversational agents. In Conversational Informatics: An Engineering Approach (pp. 49–68). Chichester: John Wiley & Sons.

Dale, R. (2008). The possibility of a pluralist cognitive science. *Journal of Experimental & Theoretical Artificial Intelligence*, 20, 155–179.

Dias, J., & Paiva, A. (2005). Feeling and Reasoning: A Computational Model for Emotional Characters. EPIA. LNCS 3808, 127-140.

Dias, J., Kiryazov, K., Paiva, A., & Grinberg, M. (2009). Implementation of emotions in TRIPLE, in preparation.

Grinberg, M., & Kostadinov, S. (2008). The Triple Model: Combining Cutting-Edge Web Technologies with a Cognitive Model in a ECA. Proc. International Workshop on Agent Based Computing V (ABC'08), Wisla, Poland

Jilka, D. J., Lebiereb, C., O'Reillyc, R. C., & Anderson, J. R. (2008). SAL: an explicitly pluralistic cognitive architecture. *Journal of Experimental & Theoretical Artificial Intelligence*, 20, 197–218.

Kabbaj, A. (2006). Development of Intelligent Systems and Multi-Agents Systems with Amine Platform. Proc. of the 14th International Conference on Conceptual Structures, ICCS 2006, 286-299.

Kiryazov, K., Petkov, G., Grinberg, M., Kokinov, B., & Balkenius, C. (2007). The Interplay of Analogy-Making with Active Vision and Motor Control in Anticipatory Robots, In Anticipatory Behavior in Adaptive Learning Systems: From Brains to Individual and Social Behavior, LNAI 4520, Springer.

Kiryazov, K. & Grinberg, M. (2009). Integrating Emotions in the Triple ECA Model, submitted to this conference.

Kokinov, B. (1994). A hybrid model of reasoning by analogy. In Advances in connectionist and neural computation theory: Vol. 2. Analogical connections (pp. 247-). Norwood, NJ: Ablex.

Kostadinov, S., & Grinberg, M. (2008). The Embodiment of a DUAL/AMBR Based Cognitive Model in the RASCALLI Multi-Agent Platform. Proc. of IVA08. LNCS 5208, pp. 356–363.

Kostadinov, S., Petkov, G., & Grinberg, M. (2008). Embodied conversational agent based on the DUAL cognitive architecture. Proc. of WEBIST 2008 International Conference on Web Information Systems and Technologies, Madeira, Portugal.

Krenn, B. (2008). RASCALLI. Responsive Artificial Situated Cognitive Agents Living and Learning on the Internet. University of Karlsruhe, Karlsruhe, Germany.

Lessmann, N., Kopp, S., & Wachsmuth, I. (2006). Situated interaction with a virtual human perception, action, and cognition. In Situated Communication (pp. 287–323). Berlin: Mouton de Gruyter.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. Psychological Review , 88, 375-407.

Ming Mao, Yefei Peng, Michael Spring (2008). Neural Network based Constraint Satisfaction in Ontology Mapping. AAAI 2008, 1207-1212

Ortony, A., Clore, G., & Collins, A. (1988). The Cognitive Structure of Emotions. Cambridge University Press, UK.

Kiryakov, A., Ognyanoff, D., & Manov, D (2005) "OWLIM – A Pragmatic Semantic Repository for OWL," Proc. Information Systems Engineering – WISE 2005 Workshops, LNCS 3807, 182-192.

Ramscar, M., & Yarlett, D. (2003). Semantic grounding in models of analogy: an environmental approach. *Cognitive Science* , 27, 41–71.

Vankov, I., Kiryazov, K., & Grinberg, M. (2008). Impact of emotions on an analogy-making robot. Proceedings of CogSci 2008, Washington DC, July 22–26