The DUAL Cognitive Architecture: A Hybrid Multi-Agent Approach

Boicho Nikolov Kokinov[1]

Abstract.[1] A hybrid (symbolic/connectionist) cognitive architecture, DUAL, is proposed. It is a multi-agent system which consist of a large number of non-cognitive, relatively simple agents, and which behaviour emerges from the behaviour of these simple agents and the interactions between them. The agents within this architecture have no internal knowledge base and goals. They are both computational devices and representational elements. They have internal (local) memory and hard-wired processes that they can run.

DUAL is hybrid at the micro level (i.e. it consists of hybrid agents) rather than at the macro level (i.e. it does not consist of separate symbolic and connectionist modules). From the symbolic perspective each agent represents a piece of world knowledge and performs some specific task while from the connectionist perspective it computes simply the activation level of the agent which reflects its relevance to the context. In this way symbolism and connectionism are considered as dualistic aspects of human cognition: the former representing the world knowledge and the latter its current relevance. The connectionist aspect of the architecture continuously "restructures" the knowledge base of the cognitive system represented by the symbolic aspect thus controlling the set of possible inferences at any moment. This makes the knowledge base dynamic and context-sensitive.

The use of the DUAL cognitive architecture in modelling similarity judgements, analogical and deductive reasoning is described.

1. INTRODUCTION

The term multi-agent systems is often used for systems of cognitive agents who interact (cooperate, negotiate, compete, contest, conflict or fight) in a common environment each of them following their own goal, i.e. these systems might be considered as models of human or animal societies. In this paper the term multi-agent system refers to a single cognitive system which consists of a large number of non-cognitive, relatively simple agents, and which behaviour emerges from the behaviour of these simple agents and the interactions between them. This concept is compatible with Minsky's Society of Mind metaphor [12].

The agents within this architecture have no internal knowledge base and goals. They are both computational devices and representational elements. They have internal (local) memory and hard-wired processes that they can run.

In the recent years interest in hybrid architectures (integrating symbolic and connectionist elements) has emerged and grown [1, 2, 4, 13, 14]. However, hybrid systems are typically designed as consisting of two or more separate modules each implementing one of the above mentioned approaches and modelling a separate cognitive process or a separate stage in a complex cognitive process. This provokes critics to argue that hybrid approaches are eclectic. In contrast, I search for an integrated architecture where symbolism and connectionism can be considered as two complementary aspects of the cognitive system and will contribute to all cognitive phenomena and all stages of the cognitive processes.

In the process of approaching this goal a hybrid cognitive architecture, DUAL, has been developed. DUAL is hybrid at the micro level (i.e. it consists of hybrid agents) rather than at the macro level (i.e. it does not consist of separate symbolic and connectionist modules). This architecture has emerged from a long-term experience in modelling various cognitive processes: memory [7], similarity judgements [9], language understanding [10], analogical reasoning [11].

2. THE DUAL COGNITIVE ARCHITECTURE: AN OVERVIEW

A cognitive system built on the DUAL cognitive architecture consists of a large number of highly interconnected hybrid agents (called DUAL agents), each of which can perform a specific task and/or represents some specific knowledge. The DUAL agents act in parallel.

Each DUAL agent consists of two parts: L-Brain and R-Brain, designed according to the symbolic and connectionist paradigm, respectively. The L-Brain of an agent represents a specific piece of real-world (or meta) knowledge, while the R-Brain represents the relevance of that knowledge to the particular context. We might think of the R-Brains as energy supplies for the corresponding L-Brains. Thus, although potentially all the L-Brains can work in parallel, in each particular moment only a small fraction of them have the necessary energy supplied by the corresponding R-Brains for actual working. The Long-Term Memory (LTM) of the cognitive system includes all the DUAL agents, while the Working Memory (WM) consists of the currently active agents. In this way the performance of the cognitive system emerges as result of the work and interaction of these currently active agents, where the set of active agents is not predefined for a specific task but is dynamic and reflects the specific context.

We might think of the symbolic and connectionist aspects of the architecture as two different viewpoints on the cognitive system using an L- or an R-filter, respectively. Thus looking through an L-filter we see only the L-Brains of the agents and perceive the cognitive system as a symbolic machine, while looking through an R-filter we see only the R-Brains of the agents and perceive the cognitive system as a connectionist machine.

The agents are linked together in a network: the whole network corresponds to the LTM of the cognitive system, whereas its active part to the WM (Figure 1). Both the agents and the links between

---

[1] Institute of Mathematics, Bulgarian Academy of Sciences, Bl.8, Acad.
G. Bonchev Street, Sofia 1113, BULGARIA,
e-mail: kokinov@bgearn.bitnet

ECAI 94. *11th European Conference on Artificial Intelligence* Edited by A. Cohn
Published in 1994 by John Wiley & Sons, Ltd.

them are interpreted differently by the symbolic and connectionist aspects of the architecture.

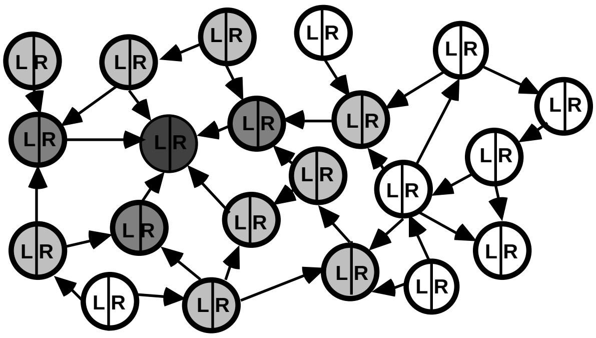

Figure 1. DUAL agents are depicted as nodes with an L and R part corresponding to their L-Brain and R-Brain correspondingly, and with filling pattern corresponding to the activation level of the agent. The network of interconnected DUAL agents corresponds to the system's LTM. The active part of the network, consisting of the agents whose activation level is above a definite threshold, correspond to the system's WM.

3. AGENTS AS REPRESENTATIONAL ELEMENTS

From the symbolic perspective the agents represent various concepts, objects, events, situations, facts, rules, plans, actions, etc. They might represent static facts as well as built-in procedural knowledge. A frame-like representation scheme is used for several reasons: 1) the integration of declarative and procedural knowledge in common structures, and 2) the possibility to have several different agents (frames) for a single object or concept reflecting different points of view. Details about the representation scheme can be found in [7]. The slot fillers are simply pointers to other frames or their slots and no special language is used for their description. This leads to a highly distributed representation of the knowledge - even a simple fact is represented by a number of interconnected agents. The links between the agents correspond to these pointers and represent various semantic links.[2]

The connectionist aspect of DUAL is used for representing context and relevance. Context is not represented by a complete explicit description of the current situation[3], i.e. by a frame within the symbolic aspect of DUAL, but instead, an implicit distributed representation is used (within the connectionist aspect) which consists of the relevance factors of each agent to the current situation. The degree of connectivity of each agent with all other agents representing elements of the current situation is chosen as a particular measure of relevance. This is called associative relevance and is represented by the activation level of the corresponding agent. Thus the activation level of the agent within the connectionist aspect represents the relevance of the description corresponding to the agent within the symbolic aspect.

The agents corresponding to entities being perceived at the moment as well as agents corresponding to the current goals of the cognitive system are called *source nodes* and they continuously emit activity, i.e. they have a constant level of activation for the period of time they are on the input or goal list. There is a relatively slow decay process so that all currently active nodes can be considered as sources of activation for a period of time. In this way the agents with a high level of activation correspond to descriptions tightly connected both with the external and internal contexts.

The links between the agents within the connectionist aspect represent the strength of the associative relations between them, i.e. how often the two agents appear in one and the same context. All the links which have some semantic interpretation within the symbolic aspect are used also by the connectionist aspect ignoring their specific semantic interpretation. In addition the *a-link*s (short for "associative links") represent arbitrary associations which are ignored by the symbolic aspect. They are not recognised by the symbolic processors and are used only by the connectionist aspect of the architecture. For example, such links are built between characteristic features often found together, between two events that have occurred within a short period of time, from class descriptions to their elements or subclasses, etc.

The links in LTM are *excitatory only*. Some symbolic processes may, however, establish additional temporary links which can be *both excitatory and inhibitory*. The weights of the links as well as the thresholds of the nodes are subject to learning. Only excitatory temporary links can become permanent.

Thus each link and each node in the network has a ***dual interpretation***: one within the symbolic representation and one within the localist connectionist network.

So, each link: 1) has a semantic label and fulfils various roles in the symbolic representation scheme, and 2) has an ascribed weight and serves to convey activation to neighbouring nodes within the connectionist network.

Each node corresponds to: 1) a *frame-like description* in the symbolic representation scheme, and 2) a simple unit in the connectionist network with an activation level corresponding to the *degree of relevance* of that conceptual description.

4. AGENTS AS PROCESSING UNITS

The R-Brains are simple connectionist processors calculating the activation values and outputs of the nodes on the basis of their input values and current activity running in parallel in a discrete synchronous manner in order to simulate the continuous process of spreading activation. They have memory for their activation value as well as for all outgoing links (pointers and weights) and a simple numeric processor working under the rules described below.

At each moment *t* every node has some activity $a_i(t)$ and passes some output - $o_i(t)$ - to its neighbours:

---

[2] The *is-a* and *instance-of* links define the concept as a specialization of or as a particular instance of a class. Each two agents linked to each other by a *c-coref* (short for `conceptual coreference') link represent one and the same entity in the world possibly from two different points of view. This allows for multiple descriptions of one and the same object, concept, situation, etc.

[3] both the *external situation* - the perceived part of the environment, and the *internal one* - the state of the mind of the cognitive system, i.e. the currently active goals, concepts, facts, etc.

$$o_i(t)=\begin{cases} 0 & \text{if } a_i(t) < \theta_i \\ p.a_i(t) & \text{otherwise,} \end{cases}$$

where $0<p<1$, and θ_i is the activation threshold. Each node receives activity from all of its neighbours and the total input from them is the weighted sum of all their outputs $net_i(t)=\sum_j \boldsymbol{w}_{ij} .o_j(t)$, where $\boldsymbol{w}_{ij}$ are w_{ij} normalised at the time of computation of $net_i(t)$ so that $\sum_j|\boldsymbol{w}_{ij}|=1$ (for all weights of links leaving an arbitrary node n_i).

There is a decay process as well, which exponentially decreases the activity of all nodes in WM with the exception of the focus (the most active node), the input and the goal nodes. Finally, the sum total of activity in a node is

$$sum_i(t)=\begin{cases} a_i(t)+net_i(t) & \text{if } n_i \text{ is the focus} \\ \tau a_i(t)+net_i(t) & \text{otherwise,} \end{cases}$$

where τ is the decay rate, $0<\tau<1$. The activation level of the node at the next moment of time is computed from this sum in the following way:

$$a_i(t+1)=\begin{cases} 0 & \text{if } sum_i(t)\le\theta_i \\ 1-(t_i/sum_i(t)) & \text{otherwise,} \end{cases}$$

where θ_i is the activation threshold for node n_i.

The L-Brains are specialised symbolic processors. They have memory for all outgoing links (pointers and labels) as well as for temporary markers (structures containing pointers to other, possibly non-neighbouring nodes). All L-Brains have the ability to receive and send markers and to differentiate links with different labels (e.g. to pass the markers only along links with specific label). In addition, the L-Brains of some agents are able to perform specific hard-wired programs corresponding to some possible actions of the cognitive system. Some examples of such specialised agents are the agents able to initiate a marker-passing process, the agents able to construct new agents (node constructors), the agents able to initiate a mapping between two descriptions, the agents able to establish local correspondence between two structures, etc.

Symbolic processors run in parallel in an asynchronous manner, each at its own individual speed. The cognitive system might have several agents with the same L-Brains (i.e. with the same hard-wired program). This permits the execution of several copies of the same program in parallel.

The differences between the L-Brains and R-Brains are summarised in Table 1.

Table 1. Agents as processing units.
Differences between the L-Brain and R-Brain processors.

Processor	Computation	Parallelism
L-Brain	symbol processing	asynchronous, individual speed
R-Brain	numeric computation	synchronous, instantaneous jumps

5. THE DUAL NATURE OF THE DUAL COGNITIVE ARCHITECTURE

The memory of the cognitive system is considered as a network of nodes and links which has a dual interpretation: 1) as a network of frame structures representing symbolically the knowledge about the world, and 2) as a localist connectionist network representing the relevance of the corresponding pieces of knowledge to the current context. The symbolic aspect reflects the static representational view on the memory, whereas the connectionist aspect reflects the dynamic process point of view on it (a particular pattern of activation corresponding to a particular context-sensitive state of the mind of the cognitive system), i.e. the symbolic aspect corresponds to the memory's external representative nature, whereas the connectionist aspect corresponds to its internal pragmatic nature.

Another dualistic approach to the DUAL agents is their interpretation both as elements of a representation and as processing units (in both the symbolic and connectionist aspects of the architecture).

6. INTERACTION BETWEEN THE SYMBOLIC AND CONNECTIONIST ASPECTS OF THE ARCHITECTURE

This interaction is realised through the close relations between the L- and R-Brain of a given agent.

The symbolic computations performed by the L-Brains might influence the work of the R-Brains. If, for example, as a result of the symbolic computation a particular agent is being put on the goal list, then it becomes a source node which means that its activation level is changed from external to the R-Brains sources. Another example is the increase of the activation level of a R-Brain when the corresponding L-Brain has received two different markers. On the other hand, when a particular L-Brain fails in doing its job for some reason then the activation of the corresponding R-Brain is decreased.

On the other hand the R-Brains influence the work of the corresponding L-Brains as well. In general the connectionist processor can be considered as an energy supply for the symbolic one, i.e. the higher the activation level of the connectionist processor, the more productive the symbolic processor.

Let us first consider the case when the agent represents a possible action of the cognitive system (including a mental one). If the activation level of that agent obtained by the R-Brain is above its threshold, then the L-Brain will start its work and it will run with a rate proportional to this activity. In this way a set of symbolic processes runs in parallel and with different rates at each particular moment.[4] These processes can communicate with each other through the links between them. Each processor has, however, a sensitivity threshold (i.e. the minimum level of activation that another node has to possess in order to be able to pass on the markers sent by this processor) associated with it which limits its communication abilities. This threshold can be absolute or relative and may depend on the activation level of that node. In this way only part of the nodes in LTM are available for the corresponding symbolic process.

Now, if a node represents a concept, object, or some other declarative knowledge, then the greater its activity, the more

---

[4] The parallel running of processes is simulated by a time-sharing organization, so that the specific rate of running a process is determined by the magnitude of the time slice allocated to it and corresponds to its level of activation.

processors will be able to use it. On the contrary, if a node is inactive, then it will be inaccessible for all processors.

Let us consider the following example of interaction between the L-Brains and the R-Brains of the agents while performing a marker-passing task. Each agent's L-Brain can pass the markers received to its neighbouring agents. However, the actual performance depends on the activation level computed by its R-Brain. So, inactive agents cannot pass any markers, while the active agents do pass the markers at a rate proportional to their activation level. In this way, in different contexts (different patterns of activation in the network of agents) the markers will follow different paths with different rates and consequently different crossing points will be found. In typical cases the shortest paths are found at first, but in some occasions a longer path might be more active and found at first (Figure 2).

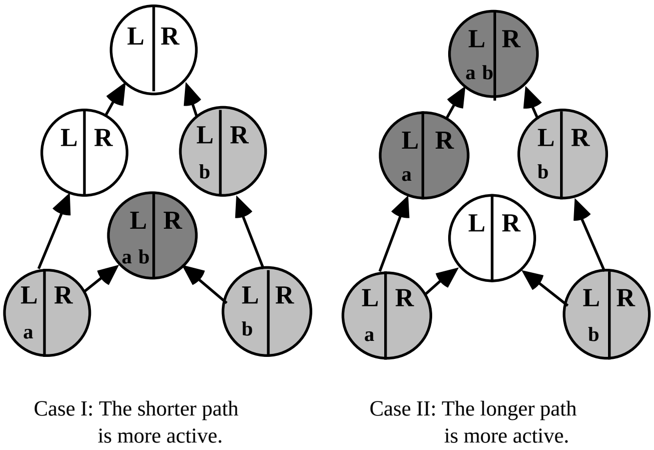

Figure 2. Context-Sensitive Marker-Passing. The level of activation of the R-Brains of the agents determines whether their L-Brains are able to pass the received markers further as well as the rate of performing of this operation. In this way the pattern of activation of the R-Brains of the agents in the network influences the emergent marker-passing process performed by their L-Brains.

In this way the connectionist aspect of the architecture continuously "restructures" the knowledge base of the cognitive system represented by the symbolic aspect thus making the knowledge base dynamic and context-sensitive.

7. EXAMPLES OF USE OF DUAL ARCHITECTURE

This architecture is implemented in COMMON LISP and simulated on an IBM-PC computer.

7.1. Semantic similarity

Let us start with a simple example: establishing the semantic similarity between two concepts. The criterion used is whether both concepts have a common superclass at any level of the hierarchy. In contrast with many other models (e.g. [6]) no restriction is made to immediate neighbours. This in combination with the possibility of having several different descriptions (possibly in different hierarchies) leads to an enormous flexibility - very far concepts might be found similar.

There are two shortcomings of such an approach: very far concepts should be considered as similar in very few occasions, and the search needed for discovering such far similarities is too demanding. Both problems are overcome in the DUAL architecture. The search even for far similarities is very efficient because only the active part of the memory is being searched, and, of course, shorter paths will be found more often.[5]

There is a side effect which makes the model of similarity judgements cognitively adequate. Computed in this way, similarity is context-sensitive: concepts which in one context are similar because have an immediate common superclass, in another context, where this superclass is not active at all, will be found dissimilar. This explains also the non-transitivity and non-symmetry of human judgements of similarity. For more detail as well as for models of other types of similarity see [9].

7.2. Analogical reasoning

A computational model of analogical reasoning, called AMBR [11], has been developed. The simulation system solves problems in the area of cooking and boiling water, eggs, etc. in the kitchen or in the forest. The knowledge base of the simulation program contains about 300 nodes and 4,000 links. There are about 10 situations related to water, three of which are the following: A) heating water on the plate of a cooking-stove in a pot, B) on the fire in a wooden vessel, and C) heating water by means of an immersion heater in a glass.

A simplified formulation of a target problem used in a psychological experiment is used as a test example in the simulation: *how can you heat water in a wooden vessel when you are in a forest, having only a knife, a match-box and an axe.* The problem is represented in the following way: the reasoner should look for a situation in which the water is in a wooden vessel and which will cause another situation in which the water will be hot and will still be in the wooden vessel.

The target problem described above has been presented several times to the system. It has been demonstrated that (in accordance with the psychological data) the behaviour of the system can vary with the variations of the context. For example, in typical situations the system finds an useless base for analogy and fails in solving the problem; with some priming (preactivation of the concept of immersion heater) the system finds a solution of the problem (putting a hot knife in the water); and in simulating the perception of a stone during the problem solving process (putting the concept of a stone on the input list) the system finds a different solution (using a hot stone instead of the hot knife).

In this way the simulation system has demonstrated AMBR's capability of analogical problem solving as well as the priming effects found in psychological experiments [8]. The simulation experiments have made also a prediction about the role of the

---

[5] The marker-passing mechanism is used, starting from the two concepts and passing two markers along *is-a*, *instance-of*, and *c-coref* links, until a cross-point is found (remember that only active agents can pass markers and that the more active agents run faster, i.e. markers pass only through active nodes, and they pass faster through more active nodes). The activation level of that cross-road is used as a measure for similarity in this context.

immediate environment and the perception of random things for the process of problem solving.

7.3. Deductive reasoning

Models of deductive reasoning are usually based on the assumption that all axioms and inference rules are available to the inference engine. However, with large knowledge bases, and particularly with humans, this is simply not true - only a small portion of the knowledge base is accessible at a given moment.

The DUAL architecture provides means for modelling this phenomenon preserving a high flexibility (without preliminary static partitions of the knowledge base). Representing all the axioms and inference rules by various agents, in each particular context only some agents will be active and consequently only some of the potential inferences will be possible. Moreover, the particular level of activation of the agents will assign different priorities to different possible inferences.

This will make deductive reasoning context-sensitive and flexible. This is in accordance with some experimental results [8] demonstrating priming effects on deductive reasoning as well. I have not yet obtained simulation results on deductive reasoning.

8. CONCLUSIONS

There is a similarity at a more abstract level between the DUAL architecture and Minsky's Society of Mind [12] both using a number of small agents which produce the meaningful behaviour at a higher level. However, the approaches used are quite different (DUAL being a hybrid architecture). A closer to DUAL hybrid architecture is used by Hofstadter and his group ([5], [3]) in modelling analogy-making. They, however, separate the declarative knowledge (Slipnet) from the procedural knowledge (Codelets) using different mechanisms for controlling them. The stochastic character of Codelets' behaviour is another important difference.

Models (of similarity and analogy) based on the DUAL architecture demonstrate high flexibility and variability in their behaviour thereby reflecting the dynamic context-sensitive nature of human cognition. On the other hand they demonstrate high efficiency restricting all searches to small parts of the knowledge base.

Work is being done on applying the DUAL architecture to modelling human deductive reasoning, decision making, and plan recognition.

It seems that this architecture can also be useful in modelling human perception and natural language understanding where the role of a dynamically evolving context is well known. However, except for its use in a word sense disambiguation task [10], I have never so far tried to build a working model in these fields.

ACKNOWLEDGEMENTS.

This research has been partially supported by the Bulgarian National Science Fund.

REFERENCES

[1] Barnden, J. & Pollack, J. (1991). High-Level Connectionist Models, Advances in Connectionist and Neural Computation Theory, vol. 1, Norwood, NJ: Ablex Publ. Corp.

[2] Dinsmore (1992). The Symbolic and Connectionist Paradigms: Closing the Gap. Hillsdale, NJ: Lawrence Erlbaum Associates.

[3] French, B. (1991). Tabletop. PhD Thesis. Univ. of Michigan.

[4] Hendler, J. (1989). Hybrid Systems (Symbolic/Connectionist). Special Issue of Connection Science, vol. 1(3).

[5] Hofstadter, D. & Mitchell, M. (in press). The Copycat Project: A Model of Mental Fluidity and Analogy-Making. In: Holyoak, K. & Barnden, J. (eds.) Analogical Connections, Advances in Connectionist and Neural Computation Theory, vol. 2, Norwood, NJ: Ablex Publ. Corp.

[6] Holyoak K. & Thagard P. (1989). Analogical Mapping by Constraint Satisfaction. Cognitive Science vol. 13, pp. 295-355.

[7] Kokinov, B. (1989). About Modeling Some Aspects of Human Memory. In: Klix, Streitz, Waern, Wandke (eds.) MACINTER II, Amsterdam: North-Holland.

[8] Kokinov, B. (1990). Associative Memory-Based Reasoning: Some Experimental Results. In: Proceedings of the 12th Annual Conference of the Cognitive Science Society, Hillsdale, NJ: Lawrence Erlbaum Associates.

[9] Kokinov, B. (1992). Similarity in Analogical Reasoning. In: B. du Boulay & V. Sgurev (eds.) Artificial Intelligence V. Amsterdam: Elsevier.

[10] Kokinov, B. (1993). Context-Sensitive Word-Sense Disambiguation. International Journal on Information Theories and Applications. vol. 1(10), pp. 47-50

[11] Kokinov, B. (in press). A Hybrid Model of Reasoning by Analogy. In: Holyoak, K. & Barnden, J. (eds.) Analogical Connections, Advances in Connectionist and Neural Computation Theory, vol. 2, Norwood, NJ: Ablex Publ. Corp.

[12] Minsky (1986). The Society of Mind. New York: Simon & Schuster.

[13] Sun, R.& Bookman, L. (1992). Integrating Neural and Symbolic Processes: The Cognitive Dimension. AAAI Press.

[14] Thornton, C. (1991) Special Issue on Hybrid Models. AISB Newsletter, No 78.