TOWARDS ACTIVE VISION IN THE DUAL COGNITIVE ARCHITECTURE

Adrian Nestor and Boicho Kokinov

Abstract: The paper describes an extension of the cognitive architecture DUAL with a model of visual attention and perception. The goal of this attempt is to account for the construction and the categorization of object and scene representations derived from visual stimuli in the TextWorld microdomain. Low-level parallel computations are combined with an active serial deployment of visual attention enabling the construction of abstract symbolic representations. A limited-capacity short-term visual store holding information across attention shifts forms the core of the model interfacing between the low-level representation of the stimulus and DUAL's semantic memory. The model is validated by comparing the results of a simulation with real data from an eye movement experiment with human subjects.

Keywords: active vision, cognitive architecture, eye movements, structural descriptions, visual attention.

Introduction

Traditional cognitive models typically isolate a small piece of a seemingly cognitive process and simulate it. However, in all domains of cognitive modeling as well as in the field of applied cognitive systems research there is a new trend towards integrating various pieces of cognitive processes and even whole processes. The current paper presents an attempt to integrate visual perception with higher-level cognitive processes like thinking and memory. The traditional goal of computer vision was to infer the structure of a three-dimensional world out of two-dimensional images with an emphasis on the lower levels of visual processing [Marr, 1982]. Automating object and scene recognition would be the next step towards an integrated visual system and, more generally, towards an integrated cognitive system.

If one gives proper credit to the claim that high-level perception, delivers the representations which form the raw material for thinking, reasoning or decision-making, then it seems critical to understand the process of constructing and making available such representations. The need to justify the format and the availability of input representations was one of the most powerful criticisms directed against traditional AI [Chalmers, French and Hofstadter, 1992]. On the other hand, integrating high-level vision with cognition in total disregard of lower-level visual processing is not necessarily a step forward. Going all the way from the visual input to high-level cognition may be the right approach but one which may be still hardly within the researchers' grasp today.

One of the most obvious difficulties with this approach may be the need to use expertise from very different fields of research. The design of cognitive architectures drawing on the joint efforts of many experts from different fields is an enterprise particularly suitable for dealing with such a difficulty. Recent attempts in this direction may be noticed. The cognitive architecture ACT-R [Anderson and Lebiere, 1998] has been lately enriched with a model of eye movements [Salvucci, 2001]. Another cognitive architecture, EPIC [Meyer and Kieras, 1997], was designed and centered on the perceptual processing of stimuli in different modalities. A similar attempt is presented below towards integrating the cognitive architecture DUAL [Kokinov, 1994, 1997; Petrov and Kokinov, 1999] with a model of visual attention and perception. Beyond DUAL's need to justify its input and to construct its own perceptual representations, this attempt is also motivated by the search for a more principled approach to modeling visual perception, an approach more faithful to the sequencing of stages and the organization of the human visual system.

Another difficulty familiar to any vision researcher is the complexity of the visual information available at the front end of the system. One way to cope with this complexity is to confine the range of possible stimuli to a predefined type. The use of microdomains with simple predefined stimuli and rules is the modeling counterpart of this experimental practice. The TextWorld microdomain, a microworld made up of blocks of text, is our candidate for this role.

However, the use of microdomains might not be enough to deal with this complexity. *Active vision,* a concept proposed by computer vision theorists in order to surpass the limitations of the image-based approach advocated by Marr [Marr, 1982] could be a realistic answer to this problem. Allowing vision to selectively attend and process parts or aspects of the available information instead of massively storing and processing all information present on the retina is not only a way to ignore irrelevant stimuli but also a necessity for a limited-capacity processing and memory system. Thus, rather than making and working on an internal copy of the outer world, an active vision system will tend to use the 'world as its own memory' [O'Regan, 1992] accessed according to the needs and the goals of the system. The serial deployment of attention and its visible counterpart, eye movements, is the way humans instantiate this principle. The model described below embodies the idea of active vision by modeling attention shifts and conditioning high-level processing of a stimulus by the availability of attention. Additionally, this offers the possibility to compare directly the performance of the model against eye movement data obtained from human subjects.

DUAL and Visual Processing

DUAL is a general cognitive architecture designed to provide a basis for modeling high-level context-sensitive cognitive processes. Although accounting for perceptual processes did not form a part of the initial motivation for its construction, a series of features characteristic to this architecture make it suitable and challenging as a framework in which to cast a model of visual perception and attention.

Hybridity in the form of a mixture of symbolic and neural network mechanisms and computations is perhaps DUAL's most significant feature. At the lowest level DUAL may be described as a large collection of units, called *DUAL agents,* reminiscent of Minsky's [Minsky, 1986] 'society of mind'. Each of these agents can be described by the symbol it stands for and by its level of activation. Agents communicate with each other both by sending symbolic messages and by spreading activation via weighted links. Coalitions of agents representing events, situations or objects tend to form themselves dynamically based on the level of activation and the links connecting a given set of agents. Finally, the set of all active units at a time may be described as the working memory of the system while the set of all units forms DUAL's long-term memory.

While hybridity has been successfully used in modeling high level-cognitive processes like analogical reasoning [Kokinov and Petrov, 1997] one can hardly imagine a domain which is more in need of such hybrid resources than visual processing. On the one hand, most of the image-based processes, e.g. the computation of a salience map [Itti and Koch, 2000], seem to be most appropriately described as the result of massively parallel numerical computations. On the other hand, perceptual primitives [Marr, 1982] and the classical structural description approach to object recognition [Biederman, 1987; Marr and Nishihara, 1978] encourage the appeal to symbolic representations and computations. In an attempt to connect a raw visual input with DUAL's semantic memory we take advantage of this hybridity by combining massively parallel activation-based computations with a serial attention-based symbolic processing mechanism instantiating the principle of active vision.

Another related point regards the way the relationship between symbolic and numerical processing is conceived of. Embracing the idea of a symbolic processor as an engine running on connectionist energy, DUAL conditions the possibility and the speed of symbolic computations on the activation level of its agents [Petrov and Kokinov, 1999]. Thus, a critical aspect in the functioning of the system is the identification of an initial source of energy or activation which enables the system to start and keep running. The current model elaborates on this topic by generalizing the concept of visual attention as a resource to be allocated [6]. Visual attention in our account is a limited and carefully managed source of energy selectively allocated to some part of the available information enabling its detailed symbolic processing. Moreover, attention spreads activation in the entire system playing the role of DUAL's energy source.

A feature which distinguishes DUAL from other architectures, typically production-rule architectures like ACT-R [Anderson and Lebiere, 1998] or EPIC [Meyer and Kieras, 1997], is the lack of a central mechanism controlling the functioning of the entire system. In DUAL each agent runs independently and in parallel with other agents using only local information obtained from its immediate neighbors. Elaborating a visual processing mechanism in the frame of a decentralized system like DUAL is surely a challenging task and a new manner of approaching vision in a cognitive architecture.

Finally, one of the most important principles underlying DUAL's development is the search for a less modular account of cognitive processes typically studied independently as part of different fields of research. Integrating memory and analogical reasoning in DUAL [Kokinov and Petrov, 1997] is one notable achievement in this direction. In the current model of visual processing we take a step forward in the same direction by describing perception as an interactive process, driven both by a low-level raw visual stimulus and by the current state and contents of DUAL's memory. Moreover, we explain and implement the categorization-based stage in object and scene recognition as a form of automatic analogy-making adding further grounds for the claim that high-level vision and analogy are at their core one and the same process [Chalmers et al., 1992].

The Model

The structure of the model is sketchily depicted in Fig 1. The visual input corresponding to a TextWorld stimulus is presented on a two-dimensional visual array representing the front end of the system. Perceptual primitives like blobs and terminations are immediately generated by cheap parallel computations. Attention is con trolled at a time by an object which allocates it selectively to some area of the stimulus. A detailed symbolic representation is constructed for this area which tends to fade away as attention is withdrawn from it and allocated to another one. Categorization takes place for the visual memory contents by retrieving and mapping object and scene categories from DUAL's semantic memory onto current visual memory representations. Each of these processes will be briefly described below.

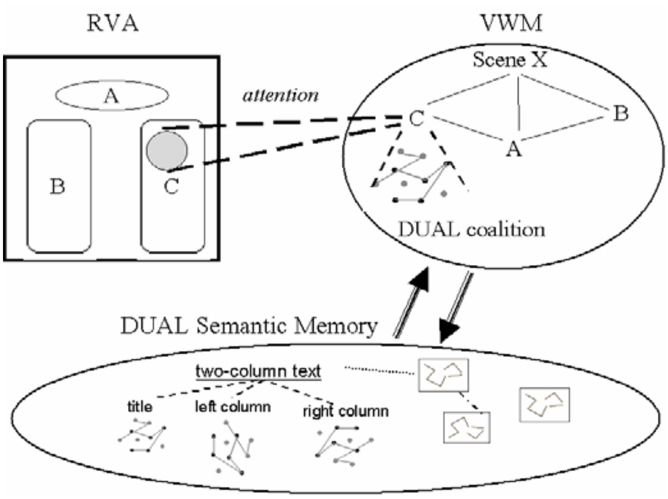

Fig. 1. The three main components of the model: the retinotopic visual array (RVA), the visual working memory (VWM) and DUAL's semantic memory. Attention is allocated to an area of the visual array by the object in VWM controlling attention while scene and object categories corresponding to the contents of VWM are retrieved from the semantic memory

TextWorld and the Retinotopic Visual Array. Stimuli are presented to the model as matrices containing filled or empty cells. Filled cells tend to group together in blocks giving the stimulus the outer appearance of a familiar configuration of typewritten text typical of TextWorld stimuli (see Fig. 2). In tests with human subjects the matrix is invisible and filled cells contain unreadable scrambled letters encouraging the subjects even more to think of the stimulus as a text format. TextWorld objects in a stimulus may also be manipulated according to specific rules giving the subjects the possibility to perform formatting tasks.

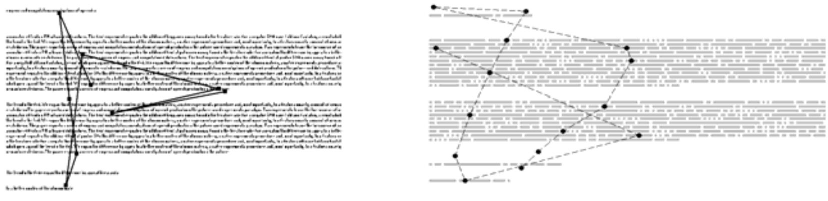

Fig. 2. A typical TextWorld stimulus with experimental (left) and simulated (right) eye movements data. Fixations are represented by points and saccades by lines connecting consecutive fixations

The type of input the model may receive for the time being is severely restricted. TextWorld stimuli are static, two-dimensional, black-and-white blocks of a roughly regular shape. The simplicity of our stimuli is the price for attempting to model vision and attention as a whole rather than focusing on a particular stage or subprocess as most models in the field do. Consequently, the types of representations and processes the models appeals to are designed to fit the specificity of our microdomain.

The first component of the model is the Retinotopic Visual Array (RVA) which in our account is a set of DUAL agents "sitting" in the cells of an imaginary matrix. Presenting a stimulus to the RVA comes down to clamping the state of each agent in RVA to filled or empty given the state of the corresponding cell in the stimulus. Each agent communicates in the array only with its immediate neighbors – at most 8 other agents – and it is unaware of its absolute coordinates in the array. By not deploying absolute coordinates the model faces a more challenging task and at the same time a more psychologically plausible one. Indeed, object representations held in VWM do not seem to be encoded in terms of absolute position but rather as configurations of stimuli in the world outside [Jiang, Olson and Chun, 2000]. Thus, we attempt to capture at least some aspects of the transition from the low-level retinotopically encoded information to the high-level spatiotopically encoded one, a transition without which visual information would be of very limited use.

Cheap parallel image-based processes take place at this level running on residual activation originating from previous or concurrent attentional processes. One such process determines whether an agent is a termination (end of row or column). For this purpose each active agent runs a local procedure checking the status of its neighbors: filled or empty. The procedure is complicated by the fact that empty cells are interspersed among filled ones just as blanks separate words in a written text. For this reason, termination detection is not always perfect although it should be reliable enough for approximating contours and edges.

A second process in the same category is blob detection. As each agent contains only local information, the perception of a cloud of filled cells as a single entity is not a trivial task. The model implements blob detection as a stepwise procedure in which neighboring cells gather initially into small groups, then intersecting groups fuse into larger ones and the process repeats itself until a whole block of cells is recognized as a single unit leading to its storing in VWM as a preattentive object [Wolfe and Bennett, 1997].

However, although cheap, the detection of perceptual primitives requires agents running symbolic procedures and is, therefore, dependent on the activation level. For this purpose, the activation spread in the system has to be a very efficient and reliable mechanism. Each DUAL agent in the system computes its activation level based primarily on the net input from its neighbors. A continuous version of the Grossberg activation function [Grossberg, 1978] governs the activation level of the agents as described below:

$$a(t_0) = a_0$$
$$da/dt= -d.a(t) + E.n(t).[M-a(t)] \qquad (1)$$

where a is the activation level as a function of time, n is the net input to the node, M is the maximal activation value, and d and E are parameters that control the rate of decay and excitation. Activation values range between 0 and M and at the beginning of each run the activation of an agent is set to a small random value. The net input is a weighted sum of the activation of neighboring agents with normalized incoming link weights. In addition, all activation levels falling below a threshold θ are set to 0. Thus, some agents in the system will be unable to run their symbolic procedures at a time at least until they recover from their shortage of energy with the aid of their neighbors.

Visual Working Memory. An impressive amount of recent data document the existence and the properties of a specialized short-term memory store holding visual information. It has been argued, for instance, that this store is different from conceptual working memory and its functioning does not interfere with the functioning of the latter [Luck and Vogel, 1997], that its organization is based on the configuration of the visual stimuli and not on their absolute locations [Jiang et al., 2000], that detailed object representation it hosts tend to fade away as attention is withdrawn from them [O'Regan, 1992] and that its capacity is of about 3-4 items [Luck and Vogel, 1997]. More importantly, it seems to encode information not in a sensorial form, i.e. complete, metrical and noncategorical, but rather in a sparse abstract code [Carlson-Radvansky and Irwin, 1995] which may recommend the use of structural descriptions as a visual representational format.

Our proposal for a Visual Working Memory (VWM) system is in line with such evidence although adapted to the peculiarities of our TextWorld stimuli. VWM is the key component in the model by its ability to construct and to hold symbolic representations of scenes and objects. A scene layout is represented in VWM as a set of objects together with their spatial relations. Similarly, an object is represented as a set of lines or bars and their spatial relations. However, the mechanisms responsible for constructing scene layouts and object representations are very different.

Object agents are generated in VWM as a result of the blob detection process. They connect with all cells which enter their composition at the RVA level as well as with other object agents via links encoding their spatial relations. Thus, a scene layout will be represented in VWM as a coalition of DUAL agents each representing an object or a relation. The detection of object relations is a preattentional process in the sense that two objects do not need to be attended in order to get connected. However, layout detection is generated at the attended area and it uses energy originating in the deployment of attention like any other symbolic process in DUAL. Markers are launched from the attention spot in different directions and each time a marker reaches RVA agents belonging to different objects these objects are connected by a link carrying the signature of the marker. For instance, if a right-propagating marker is launched by an RVA agent belonging to object A and this marker reaches an agent belonging to object B by propagating on a straight line from A to B, then the two corresponding objects are connected in VWM by a “right-of” relationship indicating that B is to the right of A. However, this mechanism is not bound to find all relations in the input as markers may get stuck on their way in inactive agents. Thus, an object in VWM may be left out of the scene coalition if it is not very active at that moment. This could account for the fact that without giving attention to an object, in addition to missing a detailed representation of that object, one may not even notice the presence of the object, a phenomenon called *inattentional blindness* or *inattentional amnesia* [Wolfe, 1999].

A mixed object and location-based account of attention control is proposed for simulating attentional shifts and for generating energy in the system. At the level of VWM, objects compete for the control of attention on the basis of their activation level – more active agents are more likely to seize attention. Once an object wins, its activation level is set to M (see equation 1) and the object starts functioning as an energy generator in the system. A second type of selection takes place at the RVA level. An object in control of attention allocates it to some fragment of itself on the RVA in the form of a limited fixed-size circular area – the ‘spotlight of attention’ [Posner, 1980]. The selection here is also based on the activation level of RVA agents competing for the role of spotlight center. As termination agents are given higher E excitation rates they tend to be more powerful competitors and to capture attention more often than other agents. RVA agents in the spotlight receive temporarily high E excitation values so that activation builds up shortly in the area. This activation is spread in turn to the RVA agents around leading to a halo of activation around the spotlight. A series of shifts on the surface of the same object is sooner or later terminated as the VWM object loses attention and another one seizes it.

Most importantly, activation enables RVA agents to support the cost of building symbolic representations. Thus, agents sufficiently active tend to group themselves in lines or bars and to generate line agents in VWM which subordinate them. Line agents form a coalition by connecting in turn the object they belong to and by connecting with each other via links encoding their spatial relations similar to the way objects connect with each other. However, unlike scene coalitions, object coalitions have only an ephemeral existence and the agents they are made up of disappear as attention is withdrawn from the object fragment they represent and their activation drops below the threshold. In this manner the model instantiates the so-called principle of *visual transience* [8]: detailed object representations fade away as attention shifts away from them.

DUAL’s semantic memory and categorization as analogy-making. The representational format utilized by DUAL’s memory and the process of analogy –making in the AMBR model have been explained and detailed elsewhere [Kokinov and Petrov, 1997] so they will not be detailed here. Our view of categorization as a form of analogy-making is based on the fact that both types of processes require the retrieval of a base from memory and its mapping with a given target. In the case of categorization the base should be a category retrieved from the semantic memory of the system and the target an instance of this category constructed on the spot.

Scene layouts in VWM are mapped onto scene categories in DUAL’s semantic memory and object representations in VWM are mapped onto object categories in the semantic memory. As category representations similarly to scene layouts and object representations in VWM are represented by DUAL coalitions, categorization is modeled by the process of mapping of different symbolic structures. Scene-context effects on object recognition can easily be explained in this framework as scene categories activate object categories which enter their composition and, therefore, facilitate the recognition of objects belonging to these categories. The top-down control of attention also falls naturally out of this schema: categories for objects which are deemed important for the current goals of the system will be represented by active agents in the semantic memory; these categories will activate their instances in VWM; finally, more active objects in VWM will seize more easily attention and will use it longer.

However, unlike analogy-making studied as a reasoning process, object and scene recognition are fast, reliable and automatic. Therefore, we view categorization as an automatic form of analogy-making or, complementarily, we explain analogy-making as an extension of categorization in the area of thinking and reasoning.

Eye movements simulation. The model has been tested by comparing its performance with the performance of human subjects in order to check its adequacy as a cognitive model. One such test is the simulation of eye movement data in an observation task with TextWorld stimuli.

Eye movements data were collected from a group of 12 subjects asked to look at a series of TextWorld stimuli including the one in Fig 2. Each stimulus was presented for 5 seconds and the resulting sequences of fixations were used for computing a transition frequency matrix for each stimulus apart. Such a matrix records the frequency of consecutive fixations in two areas of a stimulus. For instance, a cell in row A and column B represents the frequency of transitions from A to B.

The experimental data were simulated by the model as stimulus-driven overt attention shifts. For this purpose, only the performance of RVA and VWM coupled together was considered. The stimulus in Fig 2 was presented on RVA and the model was run 12 times on this stimulus while initializing RVA agents to small random values at the beginning of each run. Different sequences of attention shifts were recorded on each run mainly because of the stochastic capture of attention by objects in VWM and by object areas in RVA. A simulation run ended when the number of shifts equaled the average number of saccades for a stimulus in the experiment.

The experimental and the simulation transition frequency matrices for the stimulus in Fig 2 were not different from each other as estimated by a chi-square test ($\chi^2(24) = 23.6$, $p< 0.5$). In order to ensure this is the result of meaningful shared structure rather than the result of a lack of structure, the simulation and the experimental matrices were compared with transition frequency matrices generated from vectors recording independent fixation frequencies for each object in the stimulus [Stark and Ellis, 1981]. Both the experimental and the simulation data proved significantly different from the latter ones ($\chi^2(24) = 81.2$, $p< 0.01$; $\chi^2(24) = 76.3$, $p< 0.01$) certifying the presence of structure in the fixation sequences. However, a larger range of stimuli should certainly be tested before being able to claim that model performance and human performance are indistinguishable as far as fixation sequences are concerned.

Conclusions

The construction of a model of visual perception and attention in the framework of the cognitive architecture DUAL is advantageous both for the existing architecture and for the new model we presented above. DUAL becomes capable of processing its visual input – at least as far as TextWorld stimuli are concerned – instead of running on ready-made symbolic representations. The model, on the other hand, draws on DUAL's knowledge representation and processing mechanisms and, furthermore, it earns the ability to interface higher-level cognition. No doubt, the model should be evaluated in its own right as an artificial system whose construction is inspired by the structure and the functioning of the human visual system. For this purpose, eye movements experimental data have been compared with simulation data produced by the model alone without plugging it into DUAL's memory and resources. However, the main thrust of the model is to allow the exploration of processes emerging out of the interaction of perception with high-level cognition, e.g. the interaction of stimulus-driven and goal-directed attention control or scene context effects on object recognition. The integration of perception and cognition is a major goal for any cognitive architecture. The results above represent DUAL's first steps towards reaching this goal.

Bibliography

[Anderson and Lebiere, 1998] J. R. Anderson and C. Lebiere. The atomic components of thought. Lawrence Erlbaum Associates, Hillsdale NJ, 1998.

[Ballard, 1991] D. H. Ballard. Animate vision. Artificial Intelligence, 48 (1991) 57-86.

[Biederman, 1987] I. Biederman. Recognition by components. Psychological Review, 94 (1987) 115-177.

[Carlson-Radvansky and Irwin, 1995] L. A., Carlson-Radvanski and D. E. Irwin. Memory for structural information across eye movements. Journal of Experimental Psychology: Learning, Memory and Cognition, 21 (1995) 1441-1458.

[Chalmers, French and Hofstadter, 1992] D. Chalmers, R. French and D. Hofstadter. High-level perception, representation, and analogy: Journal of Experimental and Theoretical Artificial Intelligence, 4 (1992) 185-211.

[Eriksen and Yeh, 1985] C. W. Eriksen and Y. Yeh. Allocation of attention in the visual field. Journal of Experimental Psychology: Human Perception and Performance, 11 (1985) 583-597.

[Grossberg, 1978] S. Grossberg: A theory of visual coding, memory, and development. In: Formal Theories of Visual Perception. Ed. L. Leeuwenberg and J. Buffart. Wiley, NY, 1978.

[Hollingworth and Henderson, 2002] A. Hollingworth and J. M. Henderson. Accurate visual memory for previously attended objects in natural scenes. Journal of Experimental Psychology: Human Perception and Performance, 28 (2002) 113-136.

[Jiang, Olson and Chun, 2000] Y. Jiang, Y., I. R. Olson and M. M. Chun. Organization of visual-short term memory. Journal of Experimental Psychology: Learning, Memory and Cognition, 26 (2000) 683-702.

[Kokinov, 1994] B. Kokinov. The DUAL Cognitive Architecture: A Hybrid Multi-Agent Approach. In: Proceedings of ECAI'94. Ed. A. Cohn. John Wiley & Sons Ltd., London, 1994.

[Kokinov, 1997] B. Kokinov. Micro-Level Hybridization in the Cognitive Architecture DUAL.. In: Connectionist-Symbolic Integration: From Unified to Hybrid Architectures. Ed. R. Sun and F. Alexander. Lawrence Erlbaum Associates, Hilsdale NJ, 1997.

[Kokinov and Petrov, 1997] B. Kokinov and A. Petrov. Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: The Analogical Mind: Perspectives from Cognitive Science. Ed. D. Gentner, K. Holyoak and B. Kokinov. MIT Press, Cambridge MA, 2001.

[Itti and Koch, 2000] L. Itti and C. Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. Vision Research, 40 (2000) 1489-1506.

[Luck and Vogel, 1997] S. J. Luck and E. K. Vogel. The capacity of visual working memory for features and conjunctions. Nature, 309 (1997) 279-281.

[Marr, 1982] D. Marr. Vision. Freeman, NY, 1982.

[Marr and Nishihara, 1978] D. Marr and H. K. Nishihara. Representation and recognition of spatial organization of three-dimensional shapes. Proceedings of the Royal Society, London, Series B 200 (1978) 269-294.

[Meyer and Kieras, 1997] D. E. Meyer and D. E. Kieras. A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic Mechanisms. Psychological Review, 104 (1997) 3-65.

[Minsky, 1986] M. Minsky, M. The society of mind. Simon and Schuster, NY, 1986.

[O'Regan, 1992] J. K. O'Regan. Solving the 'real' mysteries of visual perception: The world as an outside memory. Canadian Journal of Psychology, 46 (1992) 461-488.

[Petrov and Kokinov, 1999] A. Petrov and B. Kokinov: Processing Symbols at Variable Speed in DUAL: Connectionist Activation as Power Supply. In: Proceedings of the 17th IJCAI. Ed. T. Dean. AAAI Press, 1999.

[Posner, 1980] M. I. Posner. Orienting of attention. Quarterly Journal of Experimental Psychology, 33 (1980), 3-25.

[Salvucci, 2001] D. Salvucci. An integrated model of eye movements and visual encoding. Cognitive Systems Research, 1 (2001), 201-220.

[Stark and Ellis, 1981] L. W. Stark and S. R. Ellis. Scanpaths revisited: cognitive models direct active looking. In: Eye Movements: Cognition and Visual Perception. Ed. D. Fisher, R. A. Monty and J. W. Senders. Lawrence Erlbaum Associates, Hillsdale NJ, 1981.

[Wolfe, 1999] J. M. Wolfe. Innattentional amnesia. In: Fleeting Memories. Ed. V. Coltheart. MIT Press, Cambridge MA, 1999.

[Wolfe and Bennett, 1997] J.M. Wolfe and S. C. Bennett. Preattentive Object Files. Vision Research, 37 (1997), 25-44.

Authors' Information

Adrian Nestor – Central and East European Center for Cognitive Science, New Bulgarian University, 21 Montevideo Str., Sofia 1618, Bulgaria; e-mail: adriannestor@students.nbu.bg

Boicho Kokinov – Central and East European Center for Cognitive Science, New Bulgarian University, 21 Montevideo Str., Sofia 1618, Bulgaria; Institute of Mathematics and Informatics, Bulgarian Academy of Sciences, Bl.8, Acad. G. Bonchev Str., Sofia 1113, Bulgaria; e-mail: bkokinov@nbu.bg