In Proceedings of ICCM - 2007- Eighth International Conference on Cognitive Modeling. 293 - 298. Oxford, UK: Taylor & Francis/Psychology Press.

The RecMap Model of Active Recognition Based on Analogical Mapping

Georgi Petkov (gpetkov@cogs.nbu.bg)
Central and East European Center for Cognitive Science,
New Bulgarian University
21 Montevideo Str., Sofia 1618, Bulgaria

Luiza Shahbazyan (ltsavak@abv.bg)
Central and East European Center for Cognitive Science,
New Bulgarian University
21 Montevideo Str., Sofia 1618, Bulgaria

Abstract

We propose an object recognition model based on analogical mapping and transfer. The objective of our model is to be able to generate and bind structural representations; and to recognize objects from a small set of primitives. The input is mapped to the associative memory and activation is spread upwards. Anticipations are generated through local mappings and transferred to be tested serially in the order of their relevance. Due to these mechanisms, our model is able to simulate phenomena such as object priming and global precedence effect. Additionally, it provides a framework for integrating visual perception and other higher-order cognitive processes.

Keywords: Recognition, Analogy-making, Anticipation, Attention, Binding.

Theoretical Framework

It is a textbook assumption that object recognition is a heuristic process. Rather than passive recipients of sensory data, we actively form hypotheses and make anticipations. Fortunately, we are provided with memory and rich environmental context, which constrain the number of plausible anticipations, thus enabling the fast and reliable recognition of virtually infinite number of objects, people, and events. Committed to these beliefs, we started developing a model of object recognition based on the cognitive architecture DUAL (Kokinov, 1994a). There are several premises fundamental to this model: (i) recognition is a heuristic process constructing structural representations, (ii) anticipations are generated by analogy, (iii) context supports (and sometimes hinders) recognition. These assumptions, as well as their empirical support and computational implementations, will be discussed in details in the following sections.

Object recognition is one of the most rigorously debated topics in the field of vision sciences. Although the issue has been approached from numerous paradigms and theoretical perspectives, the task of understanding and modeling vision is still far from being accomplished (see Peissig & Tarr, 2007 for a recent review). A central question is what are that type of representations on which recognition operates. There are two basic approaches: view-based and object-based theories and respectively, models. The basic idea behind the view-based approach is that there is a stored template for each of the previously seen visual patterns. Recognition involves matching the stimulus input to an existing template. Most of these models implement some kind of normalization process in order to reduce the number of the necessary templates (Poggio & Edelman, 1990).There are plenty of computational models committed to this paradigm and they are capable of simulating a vast range of psychological phenomena. However, they show weak performance in some domains such as class-level recognition, matching known to unknown viewing conditions as well as generalization (Tarr & Bulthoff, 1998).

In contrast, object-centered theories assume that recognition involves matching of view-point independent descriptions of spatial arrangements among parts of the object (Marr & Nishihara, 1978; Biederman, 1987). A classic example of this type is the recognition-by-components theory (RBC; Biederman, 1987). Its core premise is that recognition consists of extracting invariant structural representation of the object in terms of spatial relationships among basic shapes or components, the so-called geons, which are then matched to stored object representations. Furthermore, Biederman and Gerhardstein (1993; 1995) argue that all perceptual objects are decomposable and each object has a unique configuration of parts. Structural computational models were developed by Hummel and Biederman (1992) and Hummel and Stankiewicz (1996).

One question that is still open is how spatial inter-part relations originate. Because such relations are present at linguistic level, they can be readily available in the memory storage. So, the possibility remains that these relations are actively participating in object recognition through top-down mediation of the binding process. Nevertheless, it is not clear how this process is realized.

One plausible mechanism is analogy making, in particular mapping, which has already been granted the status of core principle for many cognitive processes (Gentner, Holyoak, & Kokinov, 2001). Several analogy-making computational models exist: COPYCAT (Hofstadter, 1984), ACME (Holyoak & Thagard, 1989), SME (Falkenhainer, Forbus, & Gentner, 1990), TABLETOP (French & Hofstadter, 1991), LISA (Hummel & Holyoak, 1997, 2003), AMBR (Kokinov & Petrov, 2001; Kokinov, 1994b) among others. Besides simulating analogy-making, these models demonstrate excellent performance in modeling a number of other cognitive processes such as memory retrieval (Forbus, Gentner, & Law, 1994), judgment (Petkov, 2006), and infant categorization (Kuehne, Gertner, & Forbus, 2000). Although the incorporation of perception and high-order cognition is not a new idea (Chalmers, French, & Hofstadter, 1992), none of the models, which pursues this endeavor, explores the possibility that the mapping of objects in two domains, based on the relational structure, which serves other cognitive processes, can readily account for the top-down influence of these relational structures on the binding of sensory information.

More specifically, we propose that the visual information is mapped to the information stored in memory. By analogy, anticipations are generated about the spatial relations between the elements at the input and how they can be bind to each other. Later, these anticipations are transferred and verified with the actual data. The implementation of this mechanism has one additional implication: the generated anticipations refer to all levels of the semantic hierarchy, so that binding and recognition of a higher level can precede the recognition of the lower level if the level of activation of these anticipations is higher. The global precedence phenomenon has been extensively documented (Navon, 1977; Kimchi, 1992), but it poses difficulty for the structural theories, which argue for part-base recognition (Biederman, 1987). Although our model is intrinsically structural, it overcomes this limitation due the specificity of its architecture. When the activation from the target spreads, it propagates upwards to the superordinate levels of the semantic network allowing anticipation formation at any level. As a result, the level of recognition in not fixed to a particular location within the hierarchy either.

Another important factor in object recognition is the available contextual information. Recognition performance is more accurate when the object is primed with consistent scene and dropped when the prime is inconsistent (Palmer, 1975). In addition, object detection is more accurate and naming is facilitated when the object appears in a consistent setting (Biederman, Mezzanotte, & Rabinowitz, 1982; Boyce & Pollatsek, 1992). Our model is context sensitive because the activation level of a particular bit of information basically represents its relevance to the current context. More active elements are anticipated more rigorously and verified with priority, thus speeding up their recognition

Major advantage of the proposed models is its potential for integration of visual perception with other cognitive processes such as reasoning on the basis of common mechanism, that is, mapping. Similar principle has been incorporated in COPYCAT (Hofstadter, 1994), although it focuses on higher-order perception of events and analogical reasoning. Our task is to determine the potential of mapping ability to support the cognitive system from the very beginning of visual processing, through fast, implicit recognition, to relatively slower, explicit analogy making.

RecMap Model of Active Recognition Based on Analogical Mapping

In RecMap model, the recognition involves (i) mapping of limited input information onto structurally organized memory traces; (ii) creation of anticipations on the basis of these mappings; (iii) sequential checking of the anticipations. The model is based on the cognitive architecture DUAL (Kokinov, 1994a), and builds up on the AMBR model (Kokinov & Petrov, 2001, Kokinov, 1994b) The RecMap model uses all mechanisms of AMBR and proposes new mechanisms for anticipation-forming, binding, and recognition, which are integrated with the old ones, thus allowing the mapping process to guide recognition as well.

The AMBR Model

AMBR model consists of a huge number of interacting with each other hybrid micro-agents with symbolic and connectionist part. The permanent agents (concepts and some of the instances) constitute the system's long-term memory, a semantic network with merged representation of the declarative and episodic knowledge. Each agent represents bits of information, but even small pieces of knowledge are represented by a coalition of many agents. At the same time, each agent has an activation level depending on its relevance to the ongoing context, and only active agents participate in symbolic operations.

The AMBR agents that represent the environment (source node) and the task (goal node) serve as a source of activation, which spreads with decay. Each active *instance-agent* emits a marker. This marker is sent to its parent *concept-agent* (representing type) and then upwards in the class hierarchy. When two markers meet, a *hypothesis-agent* for correspondence between the two marker-origins is created. The structural correspondence mechanism in turn creates new hypotheses on the basis of old ones. For example, if two relations are analogical, their respective arguments should also be analogical, etc. Thus, dynamically, a constraint satisfaction network of interconnected, competing with each other hypotheses emerges. Once a hypothesis maintains leading activity long enough and reaches a critical value, it is promoted to a winner, representing the analogy performed by the model.

In AMBR, the process of analogy-making is not separated into sub-processes - retrieval and mapping overlap and interact with each other. As a result, the structural constraints, crucial for the mapping process, influence the retrieval as well.

Innovations

In comparison to AMBR, RecMap is equipped with several new mechanisms. These are the creation and maintenance of hypotheses for recognition, anticipatory mechanism, and attention.

After a correspondence hypothesis emerges, a structural correspondence mechanism creates a *hypothesis for* *recognition*. For example, suppose that a certain line from the environment happens to be mapped to a particular line from memory. If the second line is a part of a square, then a recognition-hypothesis that the first line is also a part of a square emerges.

Simultaneously, the *anticipatory mechanism* (Petkov, Naydenov, Grinberg, & Kokinov, 2006) is operating. The memorized instance-agents inform the relevant relations in which they participate for all their hypotheses. If a certain relation collects the hypotheses for all its arguments, it creates an *anticipation-agent*, representing the expectation that the same relation is present in the environment. The anticipation-agents are copies of their mentor-relations but all their arguments are replaced with the respective analogical elements from the target situation.

The *attention* mechanism monitors all anticipation-agents, sorts them by their activation (i.e. relevance), and at fixed time intervals asks a simulated perceptual system to check the relation represented by the most active one. In the current version of the model the perceptual system is just a pre-defined list of the relations, which are currently present. Another role of attention is to *bind* together the hypotheses for recognition for the respective relation and their arguments. For example, because all operations are performed locally, the relational arguments may have hypothesis that are parts of a wall, but not necessarily one and the same wall. So, the binding mechanism is responsible to bind all hypothesized walls into one.

Thus, various hypotheses for recognition emerge locally, support or suppress each other, and are merged by the attentional binding mechanism. The recognition hypotheses in turn emit markers upwards in the conceptual system, participate in new hypotheses and anticipations, and create even more abstract recognition hypotheses. As a result of the relaxation of the network of hypotheses, the most active of them are promoted to winners.

Experimental Simulations

The domain for these simulations is hierarchical objects consisting of vertical or slopped lines and ovals, organized by spatial relations in figures (see Figure 1).

There are several concepts in the *long-term memory*, named 'house', 'lorry', 'tree'[1], etc., as well as their parts, and the parts of these parts and so on. For example, a particular instance-agent for a house is linked to its parts – 'roof', 'wall', and to the relation 'above' between them. In turn, the agent 'wall' is represented with a head-agent, linked to four lines, two of them horizontal, two vertical, as well as to several relations between these lines. The 'roof' consists of one horizontal, one left-right slopped, and one right-left slopped line, etc. When a target object is given to the model, it is represented only by these lines (called primitives), without any relations between them.

The model's task is to organize the square and the triangle[2] through the anticipatory and attentional mechanisms, and to create hypotheses that the square and triangle are respectively wall and roof (in competition with many other hypotheses), to anticipate possible relations between the square and the triangle, to organize them in a

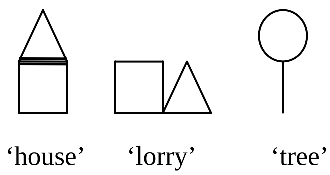

Figure 1: Example of the objects used in the simulations.

single object, and finally to recognize the house. Note that these processes overlap and the given order is very rough.

There are several instances for each concept from the *long-term memory*, and each concept is randomly linked to some of these instances, thus ensuring that the activation would propagate from the semantic (conceptual) to episodic (concrete instances) memory.

In a series of five simulations the main properties of the model are demonstrated. In the first simulation, a single object is recognized, thus the integrated work of all mechanisms is tested. In the second simulation, an object that shares all but one relation with the first one is recognized. In the third simulation, various priming effects are simulated. In the fourth and fifth simulations, the ability of the model to deal with more complex scenes and situations is tested.

Simulation 1: Recognition of a House

The task of the model in the first simulation was to recognize a single object – a 'house'. There are only seven primitives on the input, representing the straight contours – three horizontal lines, two vertical, one left-to-right slopped, and one right-to-left slopped.

The activation spreads from these primitives to the parent concepts of these lines and then back to some of their instances. Many hypotheses for correspondence emerge. For example, each of the horizontal lines creates its own hypotheses with various horizontal lines, which participate in various objects. In turn, these correspondences create hypotheses for recognition. For example, if 'line-b1' is a part of a wall, and the target 'line-1' is analogical to 'line-b1', then the respective correspondence creates and supports a hypothesis for recognition that 'line-1' is also a part of a wall.

The more instances of a particular concept are relevant, the more support the hypotheses for recognition about the

[1] All names, used in the simulations are arbitrary. The model would work as well if the agents were named ag01, ag02… The choice of the set of primitives is also arbitrary, without any psychological validation.

[2] Note, there are not any squares or triangles in the memory. The objects are organized as hypotheses for recognition as 'wall', 'cabin' (of a lorry), etc. The terms square and triangle are used only for better description.

respective concept would receive. Thus the pressure for top-down priming influence is presented. Note, however, that because of the pressure for one-to-one mapping, the hypotheses for correspondence between the target line and the stored lines inhibit each other. Thus, the contextual top-down influence is limited.

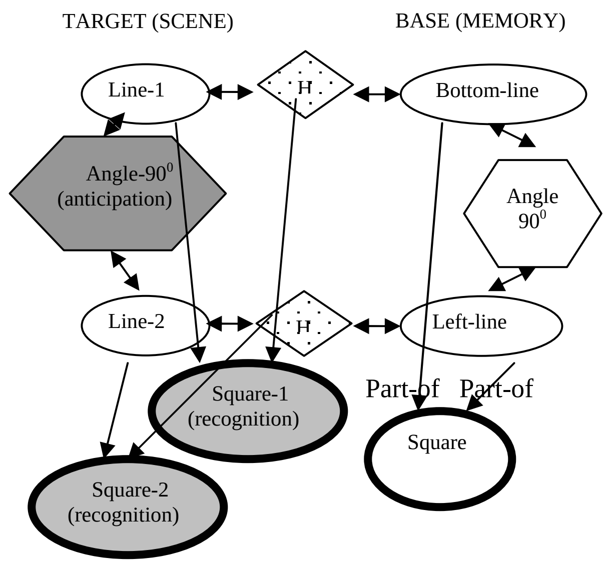

Figure 2: The hypothesis for correspondence H2 creates the recognition-hypothesis 'square-2'. Independently, H3 creates 'square-3'. In turn, the anticipation 'angle-90' is created because each of the arguments of the memorized relation 'angle-90' has hypotheses for correspondence.

At the same time, numerous anticipations about possible relations between the target lines are created (see Figure 2 for an example).

In turn, the new agents (hypotheses and anticipations) influence the spread of activation making some elements more relevant than others. The attention mechanism checks sequentially the anticipations. If certain anticipation is rejected, it just 'died'. If it is confirmed, the respective anticipation turns into an instance-agent. Thus, the description of the scene is enriched a bit. At the same time, the hypotheses for recognition of the relational arguments are bound with each other (see Figure 3). Thus, new relations are involved in the competition between the recognition-hypotheses.

Results As a result of the simulation, at time 59.24 (hundreds cycles of the program) a recognition-hypothesis 'house' becomes a winner; recognition-hypotheses for 'roof' and 'wall' become winners respectively at times 99.84 and 104.96. Interestingly, the whole object was recognized before its parts consistent with the global precedence effect demonstrated by Navon (1977). Actually, before the whole object is recognized, there is no reason to recognize the square as a wall or as a cabin for example. The resolve of the puzzle starts after confirming the anticipation 'above' with the two parts as arguments. Thus, the already created instance-agent 'above' add the decisive support for the 'house' and later on, the parts of the house are recognized as well.

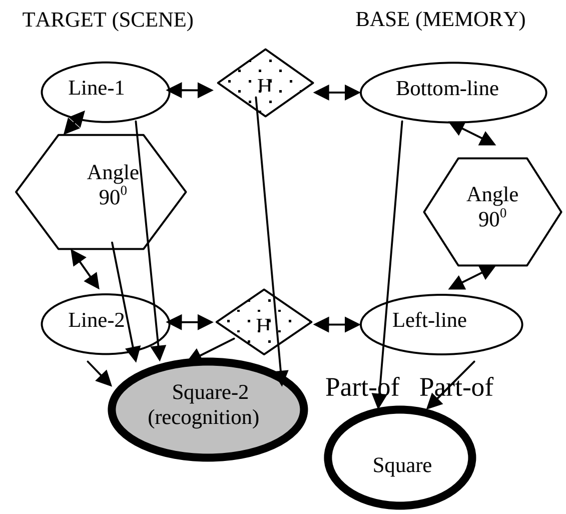

Figure 3: If the anticipation 'angle-90' is confirmed, it is transformed into an instance-agent. At the same time, the recognition-hypotheses of its arguments are bound to each other (compare to Figure 2).

Simulation 2: Recognition of a Lorry

Everything, including the seven input lines, was the same, as in the first simulation except that in the list of predefined relations the relation 'above' is replaced with 'in-touch', thus the correct response of the model was 'lorry'.

Results At time 60.20 a 'lorry' was recognized; a 'cabin' and a 'trailer' were recognized respectively at time 107.66 and 110.32.

Simulation 3: Priming

The role of the third simulation was to simulate the impact of context priming on the recognition process.. Actually, simulation 3 consists of three separate runs of the program. In all three runs, the model's task is to recognize a single lorry (the simulation 2 is repeated). However, at the first run, additional instance of a 'road', associatively linked to 'lorry' is attached to the input, thus supplying the concept 'lorry' with extra activation.

The prediction was that the higher activation of the instances of 'lorry' would facilitate the recognition process. In the second run, the concept 'fence', associated with 'house' is activated, expecting to hinder the recognition. Finally, it the third run, again 'fence' is pre-activated but with extremely high stimulation, thus simulating abnormal fixation.

Results At the first run, RecMap recognized 'lorry' at time 58.92; 'cabin' at time 86.58; 'trailer' at time 94.08, thus fully confirmed our expectations (compare with the respective results without priming from the Simulation 2 –

60.20, 107.6,6 and 110.32). At the second run the respective times were 62.22, 11.508, and 130.72. At the third run the model made wrong recognition – at time 103.12 a ‘roof’ was recognized, at time 139.28 – a ‘house, at time 178.34 – a ‘cabin’.

These results in agreement with the effects of consistent and inconsistent contextual priming demonstrated in the psychological literature (Palmer, 1975, Biederman, et al., 1982).

Simulation 4: Recognition of Two Objects

In the fourth simulation 14, instead of 7 lines were attached to the input of the model. A situation with two different objects was simulated (see Figure 4, left panel). There were not any relations between parts or primitives of different objects (i.e. anticipations for such relations were created but later rejected).

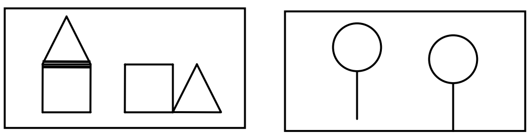

Figure 4: Left panel: Stimuli, presented to the model in simulation 4. Right panel: The base situation, used in simulation 5.

Results The overall recognition time was slowed down but not considerably. The times for recognition of a ‘lorry’, ‘cabin’, ‘trailer’, ‘house’, roof’, ‘wall’ were 58.98, 69.38, 93.04, 97.56, 122.38, 136.14, respectively. This is evidence that the model can operate on more complex scenes with a little increase of computational resources. Thus, the model’s ability to scale up is demonstrated.

Simulation 5: Integration of Recognition and Analogy-making

Finally, a whole base episode was added to the long-term memory and the capability of the RecMap model to perform the whole cycle from perception to complex analogy was tested. The base situation consists of two trees with a relation ‘left-of’ between them (see Figure 4).

Results The model successfully recognized the house and the lorry (just as in the Simulation 4) and continued with the analogy-making process.

Because all concepts for ‘house’, ‘lorry’, and ‘tree’ are sub-classes of the superordinate concept ‘neighborhood’, the respective markers from the target and base objects cross, and new hypotheses for correspondence between the target objects and the trees are created. In turn, the RecMap mechanisms created anticipations that the lorry is in left of the house, and vice versa. The former anticipation is rejected, the latter one is confirmed, and thus the right spatial analogy was settled.

Conclusions

The RecMap model for recognition, based on the DUAL architecture and the AMBR model for analogy-making is presented. The main assumptions of the model are that the analogy-making is very basic human ability, and that the recognition is an active process of dynamical creation and verification of various hypotheses and anticipations. The model is based on an associative organization of the memory; on high context sensitivity; on basic mechanisms for analogy-making and hypotheses creation; on anticipatory behavior; and on attentional mechanism for sequential testing of anticipations.

Our model was successful at stimulating several effects that are considered characteristic for human object recognition. To begin with, the model manages to anticipate, verify, and construct hierarchical structural representations of objects by analogy, which reveals the potential of this mechanism to support low as well as high-level recognition. Even more, it is able to recognize as different two objects that share all but one relation.

However, structural does not always mean part-based as we demonstrated. The recognition may start from the whole and then proceed to the parts of the objects. Even more interesting is the chronology of events when the model is presented with two objects. When the recognition began with a particular object, it continued with its parts only on the basis of the competition between the active anticipations.

Furthermore, we showed that the influence of context in the priming simulations can be modeled in an ecological manner. The model is not only able to simulate facilitative effects when the priming is consistent, but also slows down and is prone to mistakes when the context is leading.

Finally, we demonstrated that our model is able not only to bind objects and recognize them, but also to perform analogical reasoning. The novel in our approach is that all processes are guided by one and the same underlying principle - the ability of analogical mapping.

Nevertheless, there are several limitations of the current model. Although we assumed some kind of attentional mechanism, future work is needed to develop one with higher psychological plausibility. Another shortcoming is the type of the information that is used. Although the top-down construction of structural descriptions is important, it is unlikely that it is the sole pressure in recognition. Continuous metric information should be added as well as other bottom-up pressures such as salience. Finally, the model’s ability to recognize and reason about more variable and complex objects and events should be tested.

The greatest future challenge is to implement the same principles of binding on more realistic stimuli and to use both structural and metric information.

Acknowledgments

This work is supported by the Project ANALOGY: Humans – the Analogy-Making Species, financed by the FP6 NEST Programme of the European Commission. (Contr. No 029088).

References

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94*, 115–147.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 1162–1182.

Biederman, I., & Gerhardstein, P. C. (1995). Viewpoint-dependent mechanisms in visual object recognition: Reply to Tarr and Bülthoff (1995). *Journal of Experimental Psychology: Human Perception and Performance, 21*, 1506–1514.

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J.C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology, 14*, 143–177.

Boyce, S. J., & Pollatsek, A. (1992). Identification of objects in scenes: The role of scene background in object naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 531–543.

Chalmers, D. J., French, R. M., & Hofstadter, D. R. (1992). High-level perception, representation, and analogy: a critique of artificial intelligence methodology. *Journal of Experimental and Theoretical Artificial Intelligence, 4*, 185-211.

Falkenhainer, B., Forbus, K. D., and Gentner, D. (1990). The structure-mapping engine. *Artificial Intelligence, 41*, 1-63.

Forbus, K. D., Gentner, D., & Law, K. (1994). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science, 19*, 141-205.

French, R. M., & Hofstadter, D. R. (1991). Tabletop: A stochastic, emergent model of analogy-making. *Proceedings of the 13th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

Gentner, D., Holyoak, & K., Kokinov, B. (Eds.). (2001). *The analogical mind: Perspectives from cognitive science*. Cambridge, MA: MIT Press

Hofstadter, D. R. (1984). The Copycat project: An experiment in nondeterminism and creative analogies. *AI Memo 755*, Cambridge, MA: MIT Artificial Intelligence Laboratory.

Holyoak, K. J. & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99*, 480–517.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104*, 427–466.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110*, 220–264.

Hummel, J. E., & Stankiewicz, B. J. (1996). An architecture for rapid, hierarchical structural description. In T. Inui & J. McClelland (Eds.), *Attention and performance XVI: Information integration in perception and communication*. Cambridge, MA: MIT Press.

Kimchi, R. (1992). Primacy of wholistic processing and global/local paradigm: A critical review. *Psychological Bulletin, 112*, 24-38.

Kokinov, B. (1994a). The DUAL cognitive architecture: A hybrid multi-agent approach. In A. Cohn (Ed.), Proceedings *of the Eleventh European Conference on Artificial Intelligence*. London: John Wiley & Sons

Kokinov, B. (1994b). *A hybrid model of reasoning by analogy*. In K. Holyoak & J. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections*. Norwood, NJ: Ablex

Kokinov, B., & French, R. M. (2003). Computational models of analogy making. In L. Nadel (Ed.), *Encyclopedia of Cognitive Science*. London: MacMillan.

Kokinov, B., & Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In D. Gentner, K. J. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science*. Cambridge, MA: MIT Press.

Kuehne, S. E., Gentner, D. & Forbus, K. D. (2000). Modeling infant learning via symbolic structural alignment. *Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society*, 286-291.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organisation of three dimensional structure. *Proceedings of the Royal Society of London, Series B (Biological Sciences), 200*, 269–294.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9*, 353-383.

Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition*, *3*, 519-526

Peissig, J. J. & Tarr, M. J. (2007). Object recognition: Do we know more today than we did twenty years ago? Chapter for *Annual Review of Psychology ,58*, 75-96

Petkov, G. (2006). Modeling Analogy-Making, Judgment, and Choice with Same Basic Mechanisms. In D. Fum, F, Missier, & A, Stocco (Eds.), *Proceedings of the Seventh International Conference on Cognitive Modeling*. Trieste, Italy: Edizioni Goliardiche.

Petkov, G., Naydenov, Ch., Grinberg, M., & Kokinov B. (2006). Building robots with analogy-based anticipation. *In: Proceedings of the KI 2006, 29th German Conference on Artificial Intelligence*, Bremen, in press.

Poggio, T., & Edelman, S. (1990). A network that learns to recognize 3-dimensional objects. *Nature, 343,* 263–266.

Tarr, M. J., & Bülthoff, H. H. (1999). Image-based object recognition in man, monkey and machine. In M. J. Tarr & H. H. Bülthoff, *Object recognition in man, monkey and machine*. Cambridge, MA: MIT Press.