GROUNDING RELATIONS IN ACTION

Ivan Vankov **Boicho Kokinov**
i.i.vankov@cogs.nbu.bg bkokinov@nbu.bg

Central and East European Center for Cognitive Science
Department of Cognitive Science and Psychology
New Bulgarian University, 21 Montevideo Street, Sofia 1635, Bulgaria

ABSTRACT

This paper describes an attempt to ground the meaning of relational concepts in the sensory-motor dynamics resulting from our active interaction with the world. It is suggested that relations are encoded by executing or mentally simulating actions which are constrained by the environment and the specifics of our physical bodies. The virtues of the approach are demonstrated by a computational simulation of spatial relations recognition. Finally, an empirical study which renders support for the embodied view on relations is reported.

INTRODUCTION

The symbol grounding problem is one of the most pervasive issues in cognitive science. And one of the toughest aspects of the problem is the question where the meaning of relations comes from. Relational meaning is particularly hard to catch because relations do not correspond to any physical entity. The referents of relations are not directly perceivable. At the same time, relational concepts are not necessarily abstract so one can not argue that their meaning is determined by a conceptual metaphor or by their position in a semantic network.

Earlier models of relational reasoning just disregarded the problem of the origin of relational meaning. They treated relations as symbolic propositions which were provided for free at the input of the models. More recent studies attempted to solve the problem by postulating the existence of unique features - perceptual primitives, which determine the existence of a relation. Most studies in this direction were in the area of vision and focused on spatial relations (Biederman., 1987; Hummel & Biederman 1992; Regier 1996). Although they were quite successful in finding various geometric primitives which could be used as building blocks of common spatial relations (such as ‘above’, ‘left-of’, ‘in the middle’), they did not solve the general problem since they failed to account for a variety of context effects and did not provide any clue how to deal with other kinds of relations.

A more general approach to understanding the essence of relations was offered by two models that were based on models of analogy-making – SEQL (Kuehne, Forbus, Gentner, & Quinn, 2000) and DORA (Doumas, Hummel, & Sandhofer, 2008). These models proposed that relational meaning is abstracted by comparing situations in which the relation is implicitly present. The DORA model went a step further by suggesting that the representation of a relation is basically a temporal organization of lower order propositions - relational roles. Thus a progress was made by implying that the representations of relations are dynamic and that relations cannot be adequately represented by static features, whatever they are. However none of these two analogy-based models solved the relational symbol grounding problem entirely, as they assumed the existence of unknown semantic units which are required to define a relational concept or a relation role.

The idea that time plays a crucial role in relational representations has also been exploited by the proponents of the dynamic systems account. For example, Cangelosi et al. (2005) used a recurrent Elman network to catch the dynamics of visual scenes and generate a term describing the spatial relations between participating objects. In their model relational categorization was linked to the ability to predict the development of an observed movement. In another study, Williams, Beer & Gasser (2008) showed that relational knowledge could be grounded in the sensory motor loops of a simulated cognitive agent. The simulated agent was equipped with a body letting it interact with the world and the relational representations emerged as a dynamic pattern in its behaviour while solving a relational categorization task. Although such an approach to explaining cognition does not tell much about the underlying psychological principles, it raises the important issue that the dynamic nature of relations could be only revealed by tracing the interplay of action and perception in time.

RELATIONS IN ACTION

The above presented brief overview of current research on relations highlighted two major points. The first one is that relational representations need to be dynamic. The second point is that continuous interaction with the environment is needed in order to extract relational constituents. Both observations led us to the proposal that the representations of relations should involve the execution or mental simulation of actions.

The idea that relations involve performing actions is not entirely new. One of the earlier models of analogy-making (Hofstadter, 1995) represented relations as little programs – codelets - which actively seek evidence for the existence of a relation in the environment. Another model of active recognition – RecMap (Petkov & Shahbazyan, 2007) – also involved execution of actions in order to confirm or reject anticipations about existing relations. A recent model of analogy development suggested that all relations could be viewed as transformations (Leech, Mareschal, & Cooper, 2008) and thus implicitly assumed that relations should involve action.

The role of action in the representation of relations is twofold. Firstly, the execution or simulation of a specific action by itself brings meaning to the relation. Think about relations denoted by verbs such as ‘help’, ‘feed’, ‘stop’ – the literal meanings of all of them entails some physical activity. Other relations, such as ‘support’, may also be construed into a set of possibilities for action. Consider the following example ‘*the table supports the vase*’. In order to comprehend this relation, one should know what would happen to the vase if the table is moved away. To do this, he or she could either push the table (this is what infants usually do) or mentally simulate this action and imagine the behavior of the vase. If it falls on the ground, than indeed the table has been supporting it. Luckily, when people gain enough knowledge about the world they live in, they rarely need to physically execute actions in order to verify the existence of relations. However there is growing evidence that they do simulate actions (see Glenberg, 1995, for a review). Humans (and primates) are simulating actions even when perceiving scenes, in which they are not personally involved. One of the discoveries made by cognitive science during the last decades has been the mirror neurons system (Rizzolatti & Craighero, 2004). It is believed that the mirror neurons are involved in the mental simulation of actions performed by another individual. Such a system will let us understand that *loves*(‘John’, ’Marry’) when we see John hugging or kissing or even looking as if he would hug or kiss Marry.

The second role of action is to enable the role-filler binding. The idea of dynamical role-filler binding driven by temporal synchrony dates back to Shastri (Shastri & Ajjanagadde, 1993) and was exploited by the LISA model of analogy-making (Hummel & Holyoak, 1997). In LISA, the role-filler binding is implemented by consecutively activating propositions in the driver situation and making their sub-propositions fire out of synchrony. Thus the LISA model suggests that there exist a universal mechanism for dynamic binding which is used for all kinds of relations. In our view, the temporal synchronization of role-filler items is specific to each relation and is ensured by action execution. The dynamics of action execution provides natural synchronization of the corresponding entities. Consider the *supports*('table'. 'vase') example. As already mentioned, according to our account one has to try to move the table away and observe the consequences in order to comprehend this relation. The decision to execute or simulate this action could be driven either top-down or bottom-up. If it is top-down, there must have already been an anticipation that such a relation exists. Whatever the reason is, when the execution or simulation of the action is started, it naturally brings more activation to the table as it is the object that the action is applied to. At the same time executing this action activates the corresponding role of the *supports*(supporter, supported) relational concept - 'supporter' - as it stands for the relational role, which the move-away action is applied to. When the action execution is completed the vase becomes more active as it is starting to fall and at the same the second role of relational concept – 'supported' also gains activation as it stands for the object that is supposed to fall. Thus 'table' is mapped to 'supporter' and 'vase' is mapped to 'supported' and the role-filler binding problem is solved.

The following section provides an in-depth demonstration of the above mentioned ideas.

COMPUTATIONAL STUDY: SPATIAL RELATIONS CATEGORIZATION

Goals

The goal of this computational study is to provide a detailed demonstration of how the proposed embodied view on relational meaning would work within the AMBR model of analogy-making (Kokinov & Petrov, 2001).

Spatial relations were chosen for the simulation because they are a particularly good example of relations in general: they clearly have no direct physical manifestation and the range of entities that can serve as their arguments is extremely large. In fact any material object may participate in a spatial relation. Thus spatial relations pose a problem to the existing approaches to grounding relational meaning (like DORA), which assume that there are particular attributes which describe the corresponding relational roles.

Modeling architecture – AMBR

The building blocks of AMBR are hybrid nodes (agents) which exhibit both symbolic and connectionist properties. Each node has its localist meaning (an object, relation, scene, etc), but at the same time it may be a part of the distributed representation of other nodes. Some of the nodes are marked as targets and can be mapped to other, semantically similar, nodes by means of specialized 'hypothesis' nodes. Semantic similarity is dynamically computed by marker-passing over the 'ISA' links between nodes. If a target node is mapped to more than one non-target node, the rivalling hypotheses start to inhibit each other. The hypothesis competition is finally resolved by a process of constraint satisfaction. The input to the model is a source of permanent activation which is attached to some of the nodes. The activation then spreads to the rest of the network as in a classical neural network. The nodes attached to the source of activation could be changed in the course of the time.

Simulation

The simulation is based on a scenario borrowed from the RecMap model of analogy-based recognition (Petkov & Shahbazyan, 2007). The knowledge domain of RecMap consists of two-dimensional figures with mnemonic names (Figure 1). Some of the objects are ambiguous (toaster/table and flowerpot/candlestick); some differ on a single relation (house/lorry and flowerpot/lamp); some have unique features (tree). We chose the house/lorry pair for our demonstration. Apparently the only difference between these two entities is the spatial relation between their constituting parts.

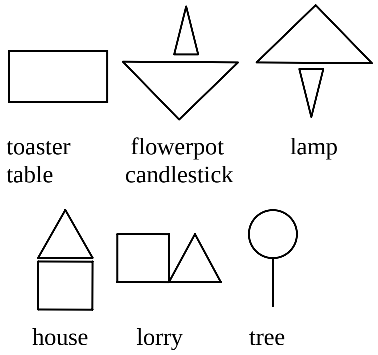

Figure 1. RecMap knowledge domain (Petkov & Shahbazyan, 2007)

Two spatial relations are crucial for this simulation – ‘above’ and ‘left-of’. In order to model their embodiment, a ‘body’ was simulated by letting the model execute four types of actions, corresponding to moving the attentional focus to four possible directions along the horizontal and the vertical axis. The representation of ‘above’ is schematically depicted in Figure 2.

Whenever a certain action is executed, the attention focus is moved in the corresponding direction. For example, when the hypothesized cognitive agents ‘looks’ upwards, the attentional focus moves to the upper object and the argument ‘upper’ of the relation ‘above’ is activated. When attention is moved downwards, the other argument ‘bottom’ is activated. The arguments themselves lack any details and have a weak link to the abstract entity ‘object’, so that any material object can map to them. The spatial relations participating in the representation of ‘house’ and ‘lorry’ are encoded by a dedicated binding node (H-ABOVE for the ‘above’ instance in ‘house’) and by linking their constituent parts to the corresponding arguments of ‘above’ and ‘left-of’.

The input to the simulation consisted of two objects – a square and a triangle. There were two items in semantic memory which were composed of a square and a triangle. That is why the only way to decide whether there was actually a house or a lorry in the environment was to take into account the spatial relations between the two perceived objects. We assumed that there is some bottom-up process that guides attention so that the agent does not look at empty places, but focuses just on the existing objects. When one of the objects is being fixated for some time, it becomes uninteresting and the probability that an action will be performed and the other object will be fixated increases.

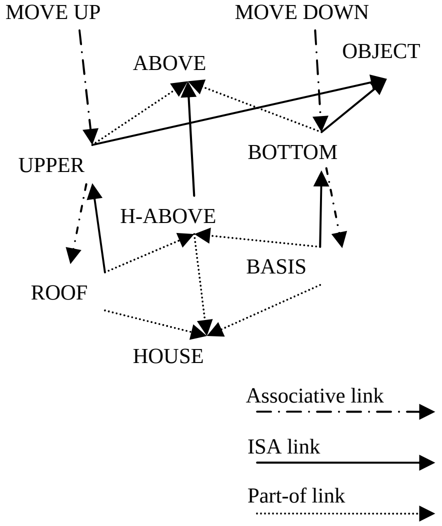

Figure 2. Representation of a spatial relation (ABOVE) and an instance of it (H-ABOVE)

An action is executed by activating the corresponding action node. The activation is not persistent but the action node retains some activation for a period time according to the decay rate of its activation function. Due to the links form the action nodes to the arguments of the spatial relations, each time the attention focus is moved by executing an action, the corresponding arguments of the relations are activated. So it happens that in approximately the same moment when the attention is redirected to a different object by executing an action, the corresponding arguments of the spatial relations are activated. The spatial relation arguments (relational roles) on their part activate those objects in memory that had been

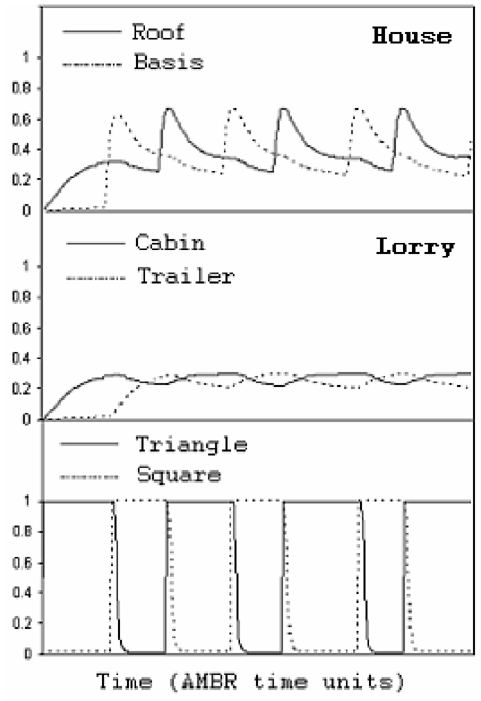

Figure 3: The dynamics of the sensory-motor loop. Note the oscillating activation levels of the two house parts which coincide in time with the changes of the perceptual input.

bound to the relational roles in the past. As a result, the activation peaks of target objects coincided in time with the activation peaks of those memorized objects which were bound to the same relation role. To put it otherwise, objects that must be mapped as corresponding arguments of the same spatial relation are activated at the same time. Thus spatial relations are actuated in the dynamics of the interaction of action and perception (Figure 3).

It turns out that ensuring that corresponding objects from memory and environment are activated at the same time is enough to bring the AMBR constraint satisfaction process to the desired outcome. This is an emergent result. It was not modelled specially for the purpose of this simulation. Figure 4 shows the outcome of the simulation with embodied relations. It is clear that when no actions were performed, the system was not able to disambiguate the perceptual input. It is notable that just a few movements were enough to resolve the constraint satisfaction network in the relevant way – a fact that is consistent with our intuition that we do not need to wonder much in order to establish a certain spatial relation.

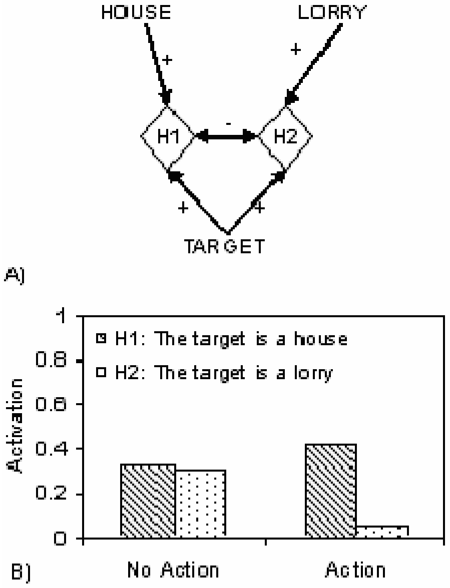

Figure 4. A) The rival recognition hypotheses characterizing the target scene. Note he inhibitory link between them. B) The activation levels of the recognition hypotheses after several movements are performed and when no action is modeled, but just the two target objects are changing turns as attentional focus. The maximum activation level is 1.0.

Conclusions

The computational study described above demonstrated how an embodied approach to representing relations would work in practice. Although the simulation was based on toy examples, its dynamics was complex enough to reveal the essence of the proposed approach – executing action dynamically binds target objects to role nodes of the relevant relations. Although the simulation was focused on the process of relational categorization, it was also demonstrated how embodied relations can be exploited in any kind of analogy-driven reasoning, as long as in the AMBR model categorization is implemented by means of the general structure mapping principles.

One issue that was not addressed in this computational study was how the proposed relational representations were acquired. Our approach of grounding relation meaning in action implies that learning relations is dependent on acquiring certain motor skills. In particular, the formation of spatial relational concepts is related to the ability to direct the focus of attention. Such a view is supported by sensory substitution studies (e.g. Collins & Bach-y-Rita), which have shown that attentional control is essential for space perception, even more important than perceptual resolution.

One of the specifics of this computational study was that it demonstrated how relations are grounded in actually (physically) executed actions. Another valid point is that the encoding of spatial relations was entirely bottom-up driven. We assume that this is the case for all relations as ubiquitous as spatial ones, when the actual execution of grounding actions comes at no or little cost. There is already evidence (e.g. Spivey & Richardson, 2001) that people tend to move their eyes when imagining spatial configurations, which renders support to our hypothesis that eye and head movements are employed in the representation of spatial relations.

Last but not least, the simulation demonstrated how the idea of grounding relations in action can be integrated into a general cognitive architecture, which aims to model cognitive phenomena ranging from low-level perception to high-level problem solving.

EMPIRICAL STUDY: EVIDENCE FOR EMBODIED FUNCTIONAL RELATIONS

The proposed approach to grounding relational concepts in action makes numerous predictions which can be tested empirically. For example, the assumption that performing or simulating action is necessary for grasping relational meaning implies that relations are dependent on the constraints of the body which executes or simulates the actions. The goal of this study is to provide evidence for such dependence.

One way to investigate the role of the body in relational representation is to exploit the idea of affordances, initially introduced by James Gibson (Gibson, 1977). Gibson defined affordances as ‘action possibilities’ - qualities of objects, or the environment, which allow an individual to perform a certain action. For example, a stairs in a house would allow a visitor to climb to the second floor, but not if the steps are 1 meter high.

Recent studies in the stimulus response compatibility paradigm have found evidence that the mere perception of an object immediately activates potential motors interactions with it (e.g. Tucker & Ellis, 1998; 2004). The specifics of these potential motors interactions are dependent on the constraints of the body of the perceiver. For example, Tucker & Ellis (1998) asked subjects to make an orientation judgement (right-side-up/upside-down) about pictures of household objects such as a coffee mug. Each object had an affordance - a handle - on the right or the left side. It was found that subjects were faster when they responded using the hand that was on the same side as the affordance.

The goal of the present study is to provide evidence that the affordances of the objects participating in a given relation influence the way the relation is processed and therefore the recognition of the relation involves simulating the action. To this end, a task was devised which required subjects to compare the relations between two pairs of objects with varying affordances.

Hypotheses

The proposed view on relations postulates that in order to compare the relations between two pairs of objects; one must perform or simulate the actions which the two relations are grounded in. Hence, the first hypothesis is that there will be an effect of the affordances of the objects on the subjects’ performance in the task.

A second prediction is that participants will solve the task more efficiently when they are able to simulate the two actions concurrently. As it was already suggested in the first part of this paper and demonstrated in the simulation study, in order to map the corresponding arguments of two instances of a relation, their activation peaks must coincide in time due to the execution of an action which participates in both representations. Hence, we expected that there would be not only a main effect of the objects' affordances, but that a specific pattern of interaction between the affordances of the objects would emerge, reflecting the constraints of the human body to execute two actions at the same time.

Method

Paricipants 48 participants (26 females, 6 left-handed) took part in the experiment for course credit or as volunteers. Their average age was 23.02 years (age range from 17 to 34, *SD* = 3.24).

Stimuli The stimulus set was constructed out of 144 photos of various household objects. Each stimulus consisted of two pairs of objects. The objects in each pair participated in a certain functional relation, such as 'hammer' – 'nail', 'key' – 'lock', 'fork' – 'spaghetti', etc. In all pairs, one of the objects had a prominent affordance – left or right. The relations in the two pairs were the same in half of the stimuli (i.e. the two pairs of objects were analogous) and different in the others (the two pairs were non-analogous). A pre-test study was used to organize the objects pairs in such a way that there was maximal agreement among people whether the relations were same or different.

All images were resized to 400x400 pixels. The two object pairs were separated by a vertical line (Figure 5). The object that had a prominent affordance was always located at the bottom.

Design The two independent variables were trial type ('analogous'/'non-analogous') and the affordance of the two object pairs. A pair of objects was made to be with left or right affordance by orienting in such a way, that it was easier to grasp the bottom object with the left or the right hand.

Thus there were four levels of the affordance factor:

LR – the affordance of the left pair was left and the affordance of the right pair was right.

RL - the affordance of the left pair was right and the affordance of the right pair was left.

LL – both affordances were left

RR – both affordances were right.

The dependent variable was the reaction time of verbal responses ('yes'/'no').

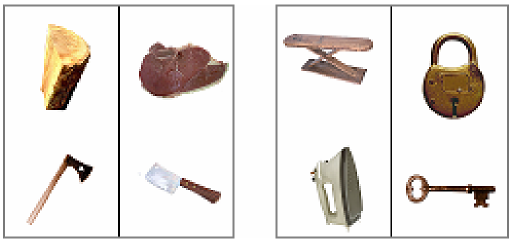

Figure 5. An example of a 'analogous'/'LR' (on the left) and a 'non-analogous'/'LL' trial (on the right).

Procedure Each stimulus was presented once to each subject. Affordance conditions and the location of objects (left/right) were counterbalanced across subjects. Analogous/non-analogous trials were randomized.

Participants were tested in a sound-proof booth. The stimuli were presented on 19" computer monitor with a resolution of 1280x1024 pixels. Before the actual experiment all participants went through a microphone training session in order to make sure that they articulate their responses clearly enough. The experimental session started with 8 practice trials, none of which appeared in the experimental part. Each trial began with a centrally location fixation cross (300ms), followed by the stimulus onset. The stimuli stayed on the screen for 5000ms or until a response subject's response was generated. Participants responded by saying 'yes' or 'no'. The intertrial interval was 2500 ms. Stimulus presentation and response recordings were controlled by E-prime software. The experimenter stayed with the subjects during the experiment and marked the response to each trial, as well as the invalid trials. The experiment took about 10 minutes.

Results

Trials in which the subject failed to respond or to articulate his or her response clearly were excluded from the analysis. A response was regarded as a correct one if the subject responded by saying ‘yes’ to a pair of relations which were designed to be analogous, and responded by ‘no’ if the trial was non-analogous. Incorrect answers were excluded from the analysis. Responses times lying more than ±2.5 standard deviations from the ‘analogous/’non-analogous’ RT means were removed as well. Thus, a total of 94.44% of the originally collected RT data were included in further analysis.

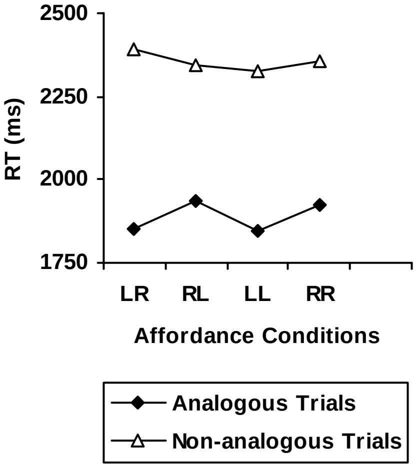

Figure 6. Experimental results. The difference between affordance conditions is significant only in the ‘same’ trials.

Analogous and non-analogous trials were analysed separately. A one way repeated measures ANOVA was performed on subject RT means and revealed significant main effect for analogous trials (*F*(3, 141) = 2.73, *p* < .05), but failed to find an effect for non-analogous trials (*F* (3, 141) = .47, *p* > .7).

A set of planned orthogonal contrast tests revealed significant difference between conditions LR and RL (*F*(3, 141) = 5.11, *p* < .05), RL and LL (*F*(3, 141) = 2.21, *p* < .05), LL and RR (*F*(3, 141) = 4.67, *p* < .05) when the two relations were analogous. There were no significant differences between affordance conditions in the non-analogous trials. The results are presented in Figure 6.

Another way to look at the data is to regard the affordances of the two pairs of objects as separate factors – ‘left affordance’ and ‘right affordance’. A 2x2 (left x right affordance) repeated measures ANOVA on subject RT means in the analogous trials revealed significant main effect of the left affordance (*F*(1, 47) = 9.99, *p* < .01) and failed to find an effect of the right one (*F*(1, 47) = .02, *p* > .8). The interaction was not significant (*F*(1, 47) = .04, *p* > .8). Similar results were found when only right-handed subjects were included in the analysis (left affordance: *F*(1, 41) = 9.80, *p* < .01; right affordance: *F*(1, 41) = .64, *p* > .4). However the pattern of results was reversed when only left-handed subject were analysed (left affordance: *F*(1, 5) = .49, *p* > .5; right affordance: *F*(1, 5) = 7.12, *p* < .05) .

Discussion

The results clearly demonstrated that comparing functional relations between objects activates motor representations related to their manipulation. Moreover, the patterns of results indicate that the effect is not due to the mere activation of possible actions. If this was the case one would expect to find a main effect of the relation which was presented near the dominant hand. Other studies have shown that the affordance effect is stronger for the dominant hand (e.g. Spivey, Richardson, & Cheung, 2001). However the result is easy to explain if we assume that people tried to simultaneously simulate the actions involved in both relations. If this is the case then the relation on the non-dominant side becomes the critical one, as subjects will have to simulate an action in an unusual way, using their non-dominant hand, while their dominant hand is engaged in another action.

In order to make sure that the observed effect is not due to automatic activation of motor representations, but to simultaneous action simulation, we ran a control study.

Control study

Participants 16 right-handed participants (11 females) took part in the study. Their average age was 23.75 years (age range from 19 to 54, *SD* = 8.36).

Stimuli The stimuli set consisted of the same objects with manipulate-able affordances which were used in the analogous trials of the previous experiment. The second objects of the corresponding relations were not displayed. Thus people saw two objects per screen (16 target trials).

Design The design involved two independent variables: the affordances of the objects located on the left and on the right side of the screen. The dependent variable was again verbal reaction time.

Procedure The setting of the experiment was the same except for the task. In this experiment, participants had to say 'yes' if none of the presented object was of natural origin and say 'no' otherwise. A set of 16 filler trials was compiled using 16 photos of man-made objects, none of which was used in the target trials, and other 16 photos of natural objects (fruits, plants, rocks, etc).

Results and Discussion Filler trials were not included in the analysis. Trials with errors as well as responses times lying more than ±2.5 standard deviations from the RT mean were also removed.

A 2x2 (left x right affordance) repeated measures ANOVA was performed on subject means reaction times. It revealed a significant main effect of the affordance of the objects located in the right side of the screen ($F(1, 15) = 4.80$, $p < .05$) and failed to find an effect of the affordance of the left objects ($F(1, 15) = .98$, $p > .3$). The interaction was not significant ($F(1, 15) = .20$, $p > .6$).

The results of the control study ruled out the possibility that the effects of the main experiment were due to automatic activation of motor representations by object perception. Thus it rendered additional support for the hypothesis that comparing relations involves the simultaneous execution or mental simulation of corresponding actions.

FINAL CONCLUSIONS

The present paper described a new approach to solving the relational symbol grounding problem by suggesting that representations of relations involve real or simulated execution of actions. A computer simulation showed that such an approach is computationally feasible and an empirical investigation managed to find evidence that it is psychologically plausible. We consider these studies as the first steps in a new line of research, which aims at revealing how our unique human ability to do complex relational reasoning is built upon the constraints of our bodies and the experience from our everyday interactions with the world.

ACKNOWLEDGEMENTS

This work was 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.

Cangelosi, A., Coventry, K., Rajapakse, R., Joyce, D., Bacon, A., Richards, L., & Newstead, S. (2005). Grounding language in perception: A connectionist model of spatial terms and vague quantifiers. In A. Cangelosi, G. Bugmann & R. Borisyuk (Eds.), Modelling language, cognitiion and action, (pp. 47-56). Singapore: World Scientific.

Collins, C., Bach-y-Rita, P. (1973). Transmission of pictorial information through the skin. Advances in Biological and Medical Physics 14, 285-315.

Doumas, L., Hummel, J., & Sandhofer, C. (2008). A theory of the discovery and predication of relational concepts. Psychological Review, 115,1-43.

Gibson, J. (1977). The Theory of Affordances. In Perceiving, Acting, and Knowing, Eds. Robert Shaw and John Bransford.

Glenberg, A. (1997). What memory is for. Behavioral and Brain Sciences, 20, 1-19.

Hofstadter, D. (1995). Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought, NY: Basic Books.

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

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

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

Kuehne, S., Forbus, K., Gentner, D., & Quinn, B. (2000). SEQL - Category learning as progressive abstraction using structure mapping. In L. Gleitman and A. Joshi (Eds.), Proceedings of the 22nd Annual Conference of the Cognitive Science Society (pp. 770-775), Mahwah, N.J: Lawrence Erlbaum.

Leech, R., Mareschal, D., & Cooper, R. (2008). Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. Behavioral and Brain Sciences, 31:4, 357-378.

Petkov, G, & Shahbazyan, L. (2007). Modelling Active Recognition as a Result of Analogical Mapping and Transfer. In D. S. McNamara & J. G. Trafton (Eds.): Proceedings of the 29th Annual Cognitive Science Society (pp. 1837-1842). Austin, TX: Cognitive Science Society.

Regier, T. (1996). The Human Semantic Potential: Spatial Language and Constrained Connectionism, Neural Network Modelling and Connectionism. Cambridge, MA: MIT Press.

Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. Annual Review in Neuroscience, 27, 169-92.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. Behavioral and Brain Sciences, 16, 417–494.

Spivey, M., & Geng, J. (2001). Oculomotor mechanisms activated by imagery and memory: eye movements to absent objects. Psychological Research, 65, 235-241.

Spivey, M., Richardson, D., & Cheung, J. (2001). Motor representations in memory and mental models: The embodied zork. Proceedings of the 23rd Annual Conference of the Cognitive Science Society (pp. 867-872), Mahwah, NJ: Lawrence Erlbaum Associates.

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. Journal of Experiemental Psychology: Human Perception and Performance, 24(3), 830-846.

Tucker M., & Ellis R. (2004). Action priming by briefly presented objects.. Acta Psychologica, 116, 185-203

Williams, P., Beer, R., & Gasser, M. (2008). An Embodied Dynamical Approach to Relational Categorization. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), Proceedings of the 30th Annual Conference of the Cognitive Science Society (pp. 223-228). Austin, TX: Cognitive Science Society.