НОВ БЪЛГАРСКИ УНИВЕРСИТЕТ

Департамент: Когнитивна наука и психология

Иван Иванов Ванков

Grounding Relations and Analogy-making in Action

А Ф Т О Р Е Ф Е Р А Т

на дисертация за присъждане на образователна и научна степен "доктор"

по професионално направление "Психология"

и научна специалност "Обща психология"

Научен ръководител:

доц. д-р Бойчо Кокинов

София, 2010 г.

Abstract

This thesis presents an attempt to ground relational concepts and relational reasoning in actually executed or mentally simulated interactions with the environment. It is suggested that relations are encoded by executing or mentally simulating actions which are constrained by the environment and the specifics of the human body. The embodiment view of relations is discussed in light of contemporary theoretical, computational and empirical research of relations and relational reasoning.

Two computer simulations based on the AMBR (Associative Memory Based Reasoning) model of analogy-making are reported. The first simulation describes how spatial relations are grounded in actually executed actions. The second simulation shows how other classes of relations, such as functional relations, are encoded and compared by simulated perceptual-motor interactions.

A series of experiments testing the predictions of the embodiment approach to relations is described. The results of the experiments indicate that people simulate actions when comparing relations and that the process of relational comparison is constrained by the characteristics of the human body.

Finally, the results of the computation and the empirical studies are put together and it is discussed to what extent the predictions of the model are supported by the experimental results. The shortcomings and limitations of the proposed approach are outlined and directions for future studies are given.

Overview

The thesis consists of 7 chapters and two appendices.

Chapter 1 introduces the problem of grounding relations and analogy-making. It is argued that plausible representations of relations are essential for adequate models of relational reasoning. A review of contemporary research reveals a converging understanding that relational representations are dynamic and interactive. It is therefore proposed that relations are grounded in action.

Chapter 2 presents the theory of embodied cognition and to relate it to the problem of grounding relational meaning. A review of empirical studies which render support for the hypothesis that relations are embodied is presented.

Chapter 3 defines the goals of the thesis.

Chapter 4 describes an extended version of the AMBR model of analogy-making (Kokinov, 1994b; Kokinov & Petrov, 2001) which is used for developing the simulations of grounded relational reasoning. A number of changes have been made to AMBR for the purpose of the current study. The representational mechanisms have been simplified and the notion of binding nodes has been introduced. A new kind of nodes has been created to represent motor knowledge. All modifications however are intact with the general principles and philosophy of AMBR.

Chapter 5 provides a more detailed description of the idea that relations are grounded in action. To this end, two computer based on the AMBR model of analogy-making are described. The goals of both simulations are to propose psychologically plausible mechanisms of embodied relational reasoning and to generate testable predictions. Simulation 1 demonstrates how spatial relations are grounded in actually executed actions. Its main message is that the representations of spatial relations involve patterns of motor programs related to attentional control. The execution of such actions does not only activate the represenations of relevant spatial relations, but also serves to solve the role-filler binding problem. The proposed mechanisms and are related to existing empirical and theoretical finding. Simulation 2 shows how another classes of relations, such as functional relations, could be grounded in simulated perceptual-motor interactions. The idea is to propose computational model which exemplifies theories of cognition by simulation (such as Barsalou, 1999; 2009; Zwaan, 2009) in the field of relational concepts and relational processing. The simulation outlines several predictions of the embodied approach to relations.

Chapter 6 describes a series of experiments which render support for our proposal that relations are grounded in the sensory-motor dynamics that characterize our interactions with the world. Experiment 1 provides initial evidence that people simulate the execution of actions when perceiving and comparing functional relations. It shows an effect of object affordances on verbal response times in a relation comparison task. A complementary experiment rules out an alternative explanation of the results. Experiment 2 extends the findings of Experiment 1 by making clear that its effects are not related to the place of presentation of the stimuli. Experiment 2 also provides additional evidence against a purely symbolic view of functional relations. Another complimentary experiment is used to show that the results of the main experiment are specific to the process of relation comparison. Experiment 3 investigates the hypothesis that an asymmetry of the human body could affect performance in a relation comparison task. It is shown that people compare functional relations faster when their bodies are manipulated symmetrically. Experiment 4 introduces the dual mental rotation task and provides further evidence for the hypothesis that relation comparison is dependent on the dynamics of perceptual-motor simulations.

Chapter 7 unifies the theoretical, computational and empiric results described in the previous chapters. We discuss whether the proposed model of grounding relations in action is able to solve the problems that motivated the present thesis. The experimental results are related to the predictions of the model. A critical review of the embodiment approach to understanding cognition is presented and the problems of embodiment are related to embodied view of relations. Finally, directions for future studies are given.

Appendix A contains the stimuli used in Experiments 1, 2 and 3.

Appendix B contains the stimuli used in Experiments 4.

Chapter 1: Introduction

The symbol grounding problem

One of the hallmarks of human cognition is the ability to use relational knowledge. It has been argued that relational reasoning lies at the core of human cognition (Hofstadter, 2000) and that it is what differentiates the human mind from the minds of other animal species (Penn, Holyoak & Povinelli, 2008).

The interest in relations has spanned a vast body of research. A number of theoretical, modelling and experimental research lines have been developed in order to elucidate the specifics of thinking about relations. There is, however, one question that has been scarcely addressed by the existing accounts of relations – how relations are represented in the mind and how they are related to the external world. In other words, how the symbol grounding problem is solved in the case of relations.

The symbol grounding problem - the problem of how symbols acquire meaning - has been a pivotal discussion point in cognitive science ever since its formulation (Harnad, 1990). After the initial boom of symbolic theories and models of cognition, it quickly became clear that no account of cognitive processes can be sound if it is unable to explain where its putative mental representations come from. The motivation for this thesis is that the same should hold for theories of relational thinking. We believe that one cannot build an adequate account of relational reasoning without a psychologically plausible model of relational representations.

The problem of grounding relational symbols is particularly difficult. Relational meaning is hard to capture because relations do not have any evident physical manifestation. There are no such material objects or events which can be pinpointed as referents of relational concepts. Relations always refer to things which cannot be directly perceived. Thus it is difficult to ground relations in physical sensations. All the same, relational concepts are not necessarily abstract in nature. Functional relations, such as 'used-for', spatial relations ('left-of') and causal relations ('is-melted-by') are among the most concrete and commonly used concepts. Therefore we cannot assume that relational meaning is always derived by abstraction from more concrete concepts.

Another problem with grounding relations is that the scope of their meaning is usually very broad. Consider the meaning of the relation 'unlocks'. Usually keys are used to unlock locks, but there are actually a countless number of other cases in which the relation between something that is locked and something else which is used to unlock could hold. For example, magnetic cards are used to unlock parking gates, remote alarm controls to unlock cars, tickets to unlock subway barriers and finger prints are used to unlock security systems. All these cases have something in common, but it is hard to define what it is. In any case it is impossible to define the meaning of 'unlocks' by a set of features describing the participating objects.

The meaning of relations is also highly context-sensitive. One can't define the relation between two objects without taking into account the whole situation in which the objects appear. The context sensitivity of relational meaning would be a serious problem to any attempt to ground relational meaning in a *fixed* set of features.

Relations and analogy-making

Analogy-making is usually defined as a process of establishing mapping between two cases or domains based on their common (relational) structure (Kokinov & French, 2002). Therefore, the problems of relational representation and relational reasoning are central to models of analogy-making.

Earlier models of analogy-making just disregarded the problem of the origin of relational meaning. For example, one of the first formal models of analogy-making - the structure-mapping theory (Gentner, 1983) and its computational instantiation (Falkenhainer, Forbus & Gentner, 1989), represented relational instances as symbolic propositions such as

‘CONTAIN(vessel, water)’. The input to the computational model was a set of symbolic structures and the reasoning processes consisted of applying rules for mapping propositions by comparing their labels. An important characteristic of the structure mapping theory is that ‘the rules depend only on the syntactic properties of the knowledge representation, and not on the specific content of the domains’ (Gentner, 1983). Hence, according to the structure-mapping theory, relational reasoning is an abstract, amodal and disembodied process, operating on structures provided by some other processes, which are not supposed to contribute to relational reasoning and therefore are not modelled.

The idea that relational reasoning should not be constrained by the nature of relational representations was also adopted by recent models of analogy-making. The ACME model (Holyoak & Thagard, 1989) represented relational symbols as localist nodes in a constraint satisfaction network. This allowed semantic and pragmatic constraints to be taken into account, but again it was assumed that the reasoning processes are not dependent on the actual meaning of relations.

The AMBR model of analogy-making (Kokinov, 1994b) introduced richer representations of relational concepts and instances by adopting a hybrid symbolic-connectionist approach. Relations were represented by computational nodes (agents) with varying processing speed, which depended on their level of relevance to the current context. Although this approach successfully accounted for a variety of context-related effects in relational reasoning (Kokinov, 1999; Kokinov & Grinberg, 2001), it did not address the question of how relations are recognized in the perceptual input and how their meaning is grounded.

The CopyCat (Mitchel & Hofstadter, 1994; Hofstadter, 1995) and TableTop (French, 1995; Hofstadter, 1995) models adopted a different view on analogy-making. Within these models, relations were represented as little programs – codelets – which actively seek evidence for the existence of an instance of a relation in the environment. This approach proved to be very successful in modelling the interplay of higher-level perception and analogy-making and suggested that the origin of relational meaning should be sought in interactions with the environment.

A more general approach to understanding the essence of relations was offered by two new models which were based on models of analogy-making – SEQL (Kuehne, Forbus, Gentner, & Quinn, 2000) and DORA (Doumas, Hummel, & Sandhofer, 2008). These models suggested that relational meaning is abstracted by comparing situations in which the relation is implicitly present. The DORA model, as well as its ancestor LISA (Hummel & Holyoak, 2003), also suggested that the representation of a relation is basically a temporal organization of lower order representations - 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, neither of these two analogy-based models solved the relational symbol grounding problem entirely, as they assumed the existence of unknown semantic units required to define a relational concept or a relational role.

A recent model of the development of analogical reasoning suggested that relations should be viewed as transformations between states (Leech, Mareschal, & Cooper, 2008). The authors also argued that there is no need to employ explicit symbols to represent relational knowledge, but it could be grounded in distributed patterns of activations in the hidden layer of a neural network. Such distributed representations of relations are unfolded in time when a relational task is to be solved.

To sum up, research so far has identified two main aspects of relations. The first one is that relational representations need to be dynamic. The second is that continuous interaction with the environment is needed in order to extract relational constituents. This thesis aims to integrate the two and put forward a new solution to the symbol grounding problem for relations, namely, that *relations are embodied*.

Embodied relations

The embodiment of relations and relational reasoning has several aspects. First, it means that the meaning of relations should be sought in potential interactions between the body of the reasoner and the environment. Such interactions could be either actually executed or simulated. The execution or simulations of actions by itself brings meaning to relations Think about relations denoted by verbs such as 'help', 'feed', 'stop' – the literal meanings of all of them entail 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 behaviour 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.

Chapter 2: The embodiment view of relations

Embodiment is a collective term encompassing a variety of approaches to understanding cognition which investigate the role of the body in shaping the mind. The embodied theories of cognition contrast traditional views which construe cognition as abstract information processing, disconnected from the outside world and independent of the physical characteristics of the system (a human body or a computer machine) which hosts the mind. Such disembodied views of cognition regard action and perception as peripheral processes which play no role in higher level cognitive functions. It is assumed that the mental representations and reasoning processes which constitute intelligence can be successfully approximated by abstract symbols and symbolic operations (Newell & Simon, 1976).

The symbolic view of cognition dominated the earlier yeas of cognitive science. There were however several lines of research which adopted a radically different stance by emphasizing the role of sensory and motor functions in higher cognition. The origins of moderns theories of embodied cognition could be traced back to, among others, the development psychology of Jean Piaget (1952), according to which human cognitive abilities emerge from sensorimotor experience; the ecological approach to visual perception of James Gibson (1979) which emphasized on the relation between perception and action; the theory of grounding abstract concepts in bodily states by means of metaphors (Lakoff & Johnson, 1980); the behaviour robotics of Rodney Brooks (1991) which demonstrated that intelligent behaviour could emerge from interactions with the world. All these lines of research headed towards the idea that human cognition is grounded in the continuous interplay of action and perception and constrained by the properties of the body and the outside world.

The accumulating evidence in support of an embodied view on cognition and the failure of the classical symbolic approach to solve a number of theoretical and practical problems led to the proliferation of new theories and models of embodied cognition (Wilson, 2002). One of the most influential ones is the theory of Perceptual Symbol Systems (PSS) proposed by Barsalou (1999). The central claim of the PSS theory is that the concepts are grounded in perceptual-motor simulations. The idea that cognition is grounded in simulation has also been exploited in language research (for reviews, see Fischer & Zwaan, 2008; Zwaan, 2009). The embodiment approach to explaining cognition has also been applied to research in memory (Glenberg, 2000; Vankov, 2009; Petkov & Nikolova, 2010; Casasanto, D., & Dijkstra; 2010), decision-making (McKinstry, Dale & Spivey 2008), mathematical reasoning (Landy & Goldstone, 2009), problem-solving (Grant & Spivey, 2003; Thomas & Lleras, 2007), emotions (Niedenthal, Barsalou, Ric & Krauth-Gruber, 2005; Oosterijk, Rotteveel, Fischer & Hess, 2009), creative cognition (Friedma & Förster, 2002), social cognition (Strack, F., Martin, L., & Stepper, 1988; Barsalou, Niedenthal, Barbey & Ruppert, 2003) and even in such abstract concepts such as goodness and badness (Casasanto, 2009).

A vast body of research has been devoted to finding empiric evidence in support of the hypothesis that conceptual knowledge is grounded in perceptual-motor interactions, but a very little part of it has been directed to relational knowledge. This is not surprising, given that relational reasoning is traditionally though to be one of the most abstract and disembodied cognitive abilities. Still, there is evidence that at least some kinds of relations are embodied.

Spatial relations are probably the most extensively studied class of relations. Most of this research has been motivated by the need to ground the meaning of spatial language. The embodiment of spatial terms and spatial relations gravitates around the idea that the apprehension of spatial relations is grounded in attention (Logan 1994, 1995; Regier & Carlson, 2001). Combined with the premotor theory of attention (Rizzolatti, 1987), according to which attention shifts utilize the same cortical circuits which are involved in action execution, the fact that spatial relations rely on attention could be considered as evidence that spatial relations are ultimately grounded in motor activity, such as eye and head movements. A number of eye-tracking studies have rendered support for the hypothesis that spatial relations are grounded in eye and head movements (Demarais & Cohen, 1998; Spivey & Geng, 2001; Coventry et al, 2009).

Another set of results supporting the idea that relations are embodied came from the studies of embodied egocentric distance and slant perception (Sinai, Ooi & He, 1999; Bhalla & Proffitt, 1999; Proffitt, Stefanucci, Banton & Epstein, 2003; Linkenauger et al. 2009). We can regard egocentric distance and slant as relational concepts as long as they involve the computation of the metric properties of a spatial relationship between two entities (e.g. self – remote object, foot of the hill-top of the hill).

The classical findings in mental imagery research can also be considered as evidence for the embodiment of relations. The connection between mental imagery and relational knowledge is in that some of the most popular tasks used to study mental imagery involved discovering an identity relationship between two objects. For example, in the seminal study of Shepard &

Metzler (1971) subjects were asked to compare two perspective drawings of objects and determined whether they depicted the same object. The main result was that the required time to recognize that two perspective drawings portray objects of the same three-dimensional shape is a linearly increasing function of the angular difference in the portrayed orientations of the two objects. Stronger evidence for the conjecture that mental transformations are implemented by perceptual-motor simulations was provided by experiments with mental rotation of body parts (Sekiyama, 1982; Parsons, 1987; Wraga, 2003; Hanakawa, Hosoda, Shindo & Honda, 2007).

Higher-order relations pose a serious problem to the theory of embodied cognition as they are highly abstract in nature and seemingly independent of any particular perceptual or motor experience. Still, Dixon & Dohn (2003) and Trudeau & Dixon (2007) managed to find evidence that the discovery of higher-order relation - alternation – is grounded in simulating the execution of actions. A similar results was reported by Day & Goldstone (2009).

To sum up, although there is a growing body of research showing that conceptual knowledge is grounded in simulated perceptual and motor activity, currently there are not many studies which address the specifics of grounding relations. The field of relational reasoning is still largely dominated by the assumptions of the symbolic view of cognition which neglects the problem of the origin of relational meaning and the nature of relational representations. The motivation of this thesis was to make an attempt to reconcile the study of relations with the emerging understanding that cognition is inherently modal, dynamic and embodied.

Chapter 3: Goals of the thesis

The goals of the thesis are twofold. First, it has to propose and describe a computational model of grounding relations and relational reasoning in the sensori-motor dynamics resulting from physically executed or simulated interactions with the environment. The model should be able to account for existing empirical findings. It should also make clear how the embodiment view of relations resolves the problems of classical approaches to relations. For example, existing models of relational reasoning assume that there exist unique perceptual or semantic features which define the meaning of relations. We have argued that in many cases it is impossible to find such features and that the meaning of relations is highly context sensitive. It was proposed that the problem of the context sensitivity of relations could be solved by grounding their meaning in actual or simulated interactions with environment. The computational model has to describe in details how relational meaning is grounded in such interactions.

One of the pivotal problems of relational reasoning is how relational arguments are bound to relational roles. It has been argued that models which lack mechanisms for role-filler binding have limited abilities to do relational reasoning (Holyoak & Hummel, 2008). The proposed computational model must include mechanisms for identifying the arguments of relations and binding them to the corresponding relational roles.

Another role of the computational model is to highlight the predictions that the proposed approach to relations makes. The model should generate testable predictions about the dynamics of embodied relational reasoning and about how the constraints of the human body could affect the discovery and comparison of relations.

The second goal of this thesis is to provide empirical evidence in support of the embodied view of relations. The results of such experiments would demonstrate that people do simulate actions when thinking about certain relations and that their performance in relational tasks is constrained by the characteristics of their bodies. It must be made certain that the results could not be explained by embodiment effects which are not specific to relational representations and relational reasoning. The experiments should also address the specific predictions which are generated by the computational model.

Chapter 4: Associative Memory Based Reasoning (AMBR)

There are several reasons why the AMBR was selected for modelling how relations are grounded in action. First, AMBR is specifically designed to deal with relations and relational reasoning and thus provides a variety of relevant tools. Other models of analogy-making, e.g. SME (Falkenheiner, Furbus & Gentner, 1989), LISA (Hummel & Holyoak, 1997), CopyCat and TableTop (Hofstadter, 1995), are also appropriate in that respect. However, the fundamental principles of AMBR - context-sensitivity, dynamicity and parallelism (Kokinov, 1998) - are particularly suited for modelling time-sensitive and interactive relational representations. Last, but not least, we have chosen to implement our model of embodied relations within AMBR in order to integrate them into a general cognitive architecture, which aims to model cognitive phenomena ranging from low-level perception to high-level problem solving.

AMBR is a model of analogy-making which is developed on top of the DUAL cognitive architecture (Kokinov, 1994a). The building blocks of AMBR are hybrid nodes (micro 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. The nodes are connected to each other by three types of connections - 'is-a', 'part-of' and 'associative'. ' Is-a' connections are used to represent conceptual hierarchy relationships, such as 'type-token' or 'class-subclass'. The role of 'part-of' links is to bring together elements which constitute a single entity, for example, all parts and properties of an object, or all elements of an event and the relationships between them Associative links are used only for spreading activation, though other type of links can also spread activation.

There are two types of special nodes in AMBR – hypotheses and binding nodes. Hypothesis nodes represent mappings (analogical connections) between other nodes. Binding nodes are used to organize nodes into coalitions which collectively represent entities such as relations, events, concepts, episodes. The same binding nodes can participate in other coalitions and server as the building block of distributed representation at a higher level.

In order to enable the representation of motor in AMBR, one more kind of special nodes was introduced: action nodes. Action nodes are used to represent motor knowledge. The difference between actions nodes and the other types of nodes lies in the activation profile of the action nodes. Once created, their activation grows for a while and then gradually fades away. The idea is that the activation of a motor program cannot be permanent. Activation of a motor program leads either to executing an action or mentally simulating it. Once the action is executed or simulated, its effect starts to decay. When the activation of an action node falls down below the working memory threshold, the action node is removed from working memory. Such a mechanism ensures that the contents of the working memory will not be cluttered with memory traces representing executed or simulated actions.

Chapter 5: Simulations

In order to show how an embodied approach to grounding relations would work we decided to model the perception of spatial relations. 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 such as DORA (Doumas, Hummer & Sandhofer, 2008) and SEQL (Kuehne, Forbus, Gentner, & Quinn, 2000), which assume that there are particular attributes which uniquely describe relational roles.

Simulation 1 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. 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). Apparently the only difference between these two entities is the spatial relation between their constituting parts. Two spatial relations are crucial for this simulation – 'above' and 'left-of'. In order to model they way they are embodied, 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 results of Simulation 1 demonstrates how an embodied approach to representing relations would work. Although the simulation was based on toy examples, its dynamics is complex enough to reveal the essence of the proposed approach – executing action binds together the corresponding arguments of different instances of the same relational concept. The simulation showed that a model of spatial relations need not rely on the existence of specific perceptual features, such as 'leftness' or 'aboveness'.

The goal of Simulation 2 was to show how functional relations can be grounded in simulated interactions with the environment. Functional relations, as well as some other classes of relations, pose a particularly hard problem to classical, disembodied, accounts of relational processing. As already discussed, the recognition of such relations is highly context sensitive and it is impossible to define a set of perceptual or semantic primitives which fully describe the existence of functional relations. Our approach to relational meaning overcomes this problem by postulating that such relations are not recognized by analyzing a static collection of the visual properties of the participating objects, but by simulating (or actually executing) potential interactions with the objects and taking into account the consequences of these interactions.

In order to model such simulated interactions we introduced a couple of new modelling tools in AMBR. The first one is the notion of affordances and the second one is *transformational* *knowledge*.

The recognition of a functional relation starts by recognizing the target objects and their spatial organization (i.e. the spatial relations between objects). The encoding of spatial relations could be a result of bottom-up driven eye or head movements as described in Simulation 1. When the initial scene is encoded, its affordances automatically activate potential motor interactions. There could be several motor programs that are activated (see Jax & Buxbaum, 2010), but the most active of them would be simulated first. The simulation of an action upon the initial scene activates relevant pieces of transformational knowledge which lead to transferring new attributes and relations to the target scene describing how it will look like if the action is actually executed. The new structures create new affordances and the process of simulating actions and predicting their consequences reiterates. At some point the description of the scene matches a certain goal state and the existence of the relation is determined.

Simulation 2 demonstrates how relations could be grounded in simulated actions. The prime role of simulated action is to discover a series of transformations of the currently available visual input which will prove the existence of a certain relation. Thus our ability to discover relations between objects is linked to the efficiency of our procedural knowledge and sensori-motor experience. Such a view of relational discovery is consistent with other computational models of relations and relation reasoning (f.e. Williams, Beer & Gasser, 2008; Cangelosi et al., 2005) as all as with the theories of embodied cognition which advocate for grounding conceptual representations and conceptual processing in perceptual-motor simulations (Barsalou, 1999; Barsalou, 2009; Zwaan, 2009).

We have found however that grounding relations in simulated actions have several other interesting implications. The simulation has raised the prediction that the performance in a relation comparison task is dependent on the ability to run two sensori-motor simulations in close temporal proximity and to align their dynamics. There is some experimental evidence in support of this hypothesis. For example, Clement (2004; 2009) reports that people use dual simulations in order to detect perceptual motor similarities between analogous situations. However there is no evidence so far that a manipulation of the ability to execute or simulate two actions simultaneously could affect performance in relation comparison task. The next chapter describes a series of experiments specially designed to test the predictions of Simulation 2.

Chapter 6: Experiments

Experimenta 1s and 1b provided initial evidence in support of the hypothesis that at least certain functional relations are embodied. First, it shows that the affordances of objects constrain performance in a relation comparison task. Second, it demonstrates that the effect of the object affordances cannot be explained by object recognition processes only.

The experiments were designed not to rely on the stimulus-response compatibility paradigm, unlike most other behavioural studies of affordances (for example Tucker & Ellis 1997, 2004; Richardson, Spivey, & Cheung, 2001; Bub & Masson, 2010). In this way it was made sure that the results could not be attributed to accidental spreading of activation from conceptual to motor areas of the brain (Mahon & Caramazza, 2008).

Experiments 2a and 2b had almost the same design as Experiment 1a and 1b. The only difference was that the stimuli were displayed one by one in the centre of the screen in order to control the order in which they were perceived and isolate the effect of presentation location. The results replicated the findings of Experiment 1a and 1b, as long as an effect of the object affordances of the relations was found. The contribution of Experiment 2a is that results cannot be explained by presentation location as all stimuli were presented in the centre of the screen. The elimination of the factor of presentation location served to set apart the effect of the affordances of the stimuli from any spatial compatibility effects.

The goal of Experiment 3 was to test in a more explicit way the hypothesis that the asymmetries of the human body may constrain the process of relation comparison. The results of the experiment are twofold. First, it was demonstrated that manipulations of subjects' bodies can have an effect in a task, which seemingly involves no limb-based motor activity. Second, we found that subjects responded faster when the manipulation of their bodies was symmetric. This result renders support for the hypothesis that the process of relation comparison is dependent on specific constraints of the human body, such as the symmetry of the limbs.

Experiment 3 demonstrated that people compare faster functional relations when their body is manipulated symmetrically. We explained this finding by assuming that asymmetric manipulations of the subjects' arms resulted in an inability to dynamically align the two perceptual-motor simulations. To further test this hypothesis, we designed and conducted Experiment 4, in which the dynamics of the simulated interactions underlying relation meaning is explicitly manipulated. The experiment was based on the well known and extensively studied mental rotation paradigm (Shepard & Metzler, 1971). The results of the experiment confirmed that people compare relations faster when they are able to dynamically align the perceptual-motor simulations underlying the relational representations.

Chapter 7: General discussion and conclusions

Contributions

In our view, the proposed approach of representing and processing relations is able to solve the problem of grounding relational symbols. We proposed that relations are grounded in simulated or physically executed interactions with the environment. The nature of these interactions could vary immensely and therefore the approach is applicable to various domains of relational knowledge. A review of existing empiric research provided converging evidence that various kinds of relations are grounded in the interplay of action and perception. Two computational studies demonstrated how relations of very different nature – spatial relations and functional relations – could be represented in terms of executed or simulated actions.

Grounding relations in action ensures that relational representations are context-sensitive. In our proposal, the encoding of a relation involves a physical or simulated interaction between the body of the reasoner and the outside world. Such an interaction would be situated in the current context, including the physical properties of the environment and bodily constraints and mental states of the reasoner. In general, the intrinsic context sensitivity of embodied representations is regarded as one of the strongest arguments in support of the embodiment approach to cognition (Barsalou, 1999).

The embodiment approach to relations generates a number of predictions which can be used to falsify it. The predictions of the embodiment approach are easier to be tested psychologically than the conjectures of traditional symbolic accounts of relational reasoning. Chapter 6 describes a series of experiments which tested some of the predictions outlined by Simulation 2.

Short-comigns and limitations

Although the current work makes progress in explaining the origin of relational representations and relational reasoning, it fails to address a number of important issues. Probably the most serious short-coming is that it does not include a model of how embodied relational representations are developed. The proposed approach of grounding relations in action suggests that the formation of relational concepts is dependent on acquiring specific motors skills and gaining sensorimotor experience. A model of learning embodied relations should be able to show how they emerge as a result of continuous interactions with the environment.

One of the motivations to use AMBR to model the grounding of relations was that such an approach would allow using embodied relation representations in various analogy-making tasks. The presented simulations however demonstrated only rudimentary forms of analogy-making with embodied relations, such as relation categorization and comparison of relations. We did not show how such representations would affect specific aspects of analogy-making, such as cross-mapping and aligning higher-order relational structures.

In our view, the execution or simulation of interactions with the environment is necessary in order to discover relations. The results of the experiments described in this thesis indicate that people run perceptual-motor simulations when comparing relations. However it is possible that the role of such simulations is only in enhancing or impeding the processing of relations.

We have proposed that grounding relations in action is a general solution to the problem of the origin of relational meaning. We decided to start the investigation of embodied relations with spatial and functional relations because there was already some evidence that they are embodied and it was easier to specify the actions underlying their meaning. However it remains an open question whether the same approach is applicable to all relations. For example, it is not clear what actions possibly underlie the meaning of relations such as ‘part-of’, ‘is-a’, ‘sister-of’ or ‘better-than’.

References

Barsalou, L. W. (1999). Perceptual symbol systems. *The Behavioral and brain sciences*, *22(*4).

Barsalou, L.W. (2009). Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London: Biological Sciences, 364*, 1281-1289.

Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, *43*, 43-92. San Diego: Academic Press.

Bhalla, M., & Proffitt, D. R. (1999). Visual-motor recalibration in geographical slant perception. *Journal of experimental psychology: Human perception and performance*, *25*(4), 1076-1096.

Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence.* , *47*, 139-159.

Bub, D. N., & Masson, M. E. J. (2010). Grasping beer mugs: On the dynamics of alignment effects induced by handled objects. *Journal of Experimental Psychology: Human Perception and Performance*, *36(2)*, 341-358.

Cangelosi, A., Coventry, K. R., Rajapakse, R., Joyce, D. A. N., Bacon, A., Richards, L. Y. N. N., & Newstead, S. N. (2005). Grounding language in perception: A connectionist model of spatial terms and vague quantifiers. In A. Cangelosi, G. Bugmann, & R. Borisyuk (Eds.) Modelling Language, Cognition and Action: Proceedings of the 9th Neural Computation and Psychology Workshop, (pp. 47-56). Singapore: World Scientific.

Casasanto, D. (2009). Embodiment of abstract concepts: Good and bad in right- and left-handers. *Journal of Experimental Psychology: General*, *138(3)*.

Casasanto, D., & Dijkstra, K. (2010). Motor action and emotional memory. *Cognition, 115(1)*.

Clement, J, (2004). Imagistic processes in analogical reasoning: Conserving transformations and dual simulations. In K. Forbus, D. Gentner, & T. Regier (Eds.), Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society (pp. 233-238). Mahwah, NJ: Erlbaum.

Clement, J. J. (2009). The role of imagistic simulation in scientific thought experiments. *Topics in Cognitive Science, 1(4)*, 686-710.

Coventry, K. R., Lynott, D., Cangelosi, A., Monrouxe, L., Joyce, D., & Richardson, D. C. (2009). Spatial language, visual attention, and perceptual simulation*. Brain and Language, 112(3),* 102-13*.*

Day, S., & Goldstone, R. (2009). Analogical transfer from interaction with a simulated physical system. In N. Taatgen, & H. van Rijn (Eds.) Proceedings of the 31th Annual Conference of the Cognitive Science Society , (pp. 1406-1411). Austin, TX: Cognitive Science Society.

Demarais, A. M., & Cohen, B. H. (1998). Evidence for image-scanning eye movements during transitive inference. *Biological psychology , 49 (3),* 229-247

Dijkstra, K., Kaschak, M. P., & Zwaan, R. A. (2007). Body posture facilitates retrieval of autobiographical memories*. Cognition, 102(1)*, 139-149.

Dixon, J. A., & Dohn, M. C. (2003). Redescription disembeds relations: evidence from relational transfer and use in problem solving. *Memory & cognition, 31(7),* 1082-1093.

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

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: algorithm and examples*. Artificial Intelligence, 41,* 1-63.

Fischer, M. H., & Zwaan, R. A. (2008). Embodied language: a review of the role of the motor system in language comprehension. *Quarterly journal of experimental psychology, 61(6),* 825-850.

French, R. (1995). Tabletop: An Emergent, Stochastic Computer Model of Analogy-Making. MIT Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy*. Cognitive Science, 7(2),* 155-170.

Gibson, J.J. (1979). The Ecological Approach to Visual Perception. Boston: Houghton Mifflin

Glenberg, A. M. (2000). What memory is for. Behavioral and Brain Sciences, 20(01), 1-19.

Grant, E. R., & Spivey, M. J. (2003). Eye movements and problem solving: Guiding attention guides thought. *Psychological Science, 14*, 462-466.

Hanakawa, T., Hosoda, C., Shindo, S., & Honda, M. (2007). Mental rotation of hands and feet involves somatotopically organized brain regions*. Neuroscience Research, 58*, S60.

Harnad, S. (1990) The Symbol Grounding Problem*. Physica D, 42*, 335-346.

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

Hofstadter, D. (2000). Analogy as the Core of Cognition. In The Analogical Mind: Perspectives from Cognitive Science, Dedre Gentner, Keith J. Holyoak, and Boicho N. Kokinov (eds.). Cambridge MA: The MIT Press/Bradford Book.

Hofstadter, D. & Mitchell, M. (1994). The Copycat project: A model of mental fluidity and analogy-making. In K. Holyoak and J. Barnden (Eds), Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections. Norwood, NJ: Ablex Publishing Corporation

Holyoak, K., & Hummel, J. E. (2008). No way to start a space program. Commentary on R. Leech, D. Mareschal and R. P. Cooper, Analogy as relational priming: A developmental and computational perspective on the origins of a complex skill, *Behavioral and Brain Sciences, 31,* 388-389.

Holyoak K., & Thagard P. (1989). Analogical Mapping by Constraint Satisfaction. *Cognitive Science, 13,* 295-355.

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

Jax, S. A., & Buxbaum, L. J. (2010). Response interference between functional and structural actions linked to the same familiar object. *Cognition* , *115* (2), 350-355.

Kokinov, B. (1994a). The DUAL cognitive architecture: A hybrid multi-agent approach. In A. Cohn (Ed.) European Conference on Artificial Intelligence, (pp. 203-207). 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 (pp. 247- 318). Norwood, NJ: Ablex.

Kokinov, B. (1998). Analogy is like cognition: dynamic, emergent, and context-sensitive In K. Holyoak, D. Gentner, & B. Kokinov (Eds.), Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences (pp. 96-105). Sofia: NBU Press.

Kokinov, B. (1999). Dynamics and Automaticity of Context: A Cognitive Modelling Approach. In: Bouquet, P., Serafini, L., Brezillon, P., Benerecetti, M., Castellani, F. (eds.) Modeling and Using Context. Lecture Notes in Computer Science (Lecture Notes in Artificial Intelligence), vol. 1688, Springer Verlag.

Kokinov, B., French, R. (2002). Computational Models of Analogy-Making. In: Lynn Nadel (ed.) Encyclopedia of Cognitive Science. Macmillan.

Kokinov, B., Grinberg, M. (2001). Simulating Context Effects in Problem Solving with AMBR. In: Akman, V., Thomason, R., Bouquet, P. (eds.) Modeling and Using Context. Lecture Notes in Computer Science (Lecture Notes in Artificial Intelligence), vol. 1775, Springer Verlag.

Kokinov, P., & Petrov, A. (2001). Integration of memory and teasoning in analogy-making: the AMBR model.. 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, & A. Joshi (Eds.) Proceedings of the 22nd Annual Conference of the Cognitive Science Society, (pp. 770-775). Mahwah, NJ: Lawrence Erlbaum.

Landy, D., & Goldstone, R. L. (2009). How much of symbolic manipulation is just symbol pushing? In N.A. Taatgen & H. van Rijn (Eds.), Proceedings of the 31th Annual Conference of the Cognitive Science Society (pp. 1318-1323). Austin, TX: Cognitive Science Society.

Lakoff, G., & Johnson, M. (1980). Metaphors we live by. Chicago, IL: University Of Chicago Press.

Leech, R., Mareschal, D., & Cooper, R. P. (2008). Analogy as relational priming: A developmental and computational perspective on the origins of a complex cognitive skill. *Behavioral and Brain Sciences, 31(04),* 357-378.

Linkenauger, S. A., Witt, J. K., Stefanucci, J. K., Bakdash, J. Z., & Proffitt, D. R. (2009). The effects of handedness and reachability on perceived distance. *Journal of experimental psycholog: Human perception and performance, 35(6),* 1649-1660.

Logan, G. D. (1994). Spatial attention and the apprehension of spatial relations. *Journal of experimental psychology: Human perception and performance, 20(5)*, 1015-1036.

Logan, G. D. (1995). Linguistic and conceptual control of visual spatial attention. *Cognitive psychology, 28(2),* 103-174.

Mahon, B., & Caramazza, A. (2008). A critical look at the embodied cognition hypothesis and a new proposal for grounding conceptual content*. Journal of Physiology-Paris, 102(1-3),* 59-70.

McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science, 19(1),* 22-4.

Newell, A, & Simon, H (1976). Computer Science as Empirical Inquiry: Symbols and Search. Communications of the ACM, 19(3), 113-126.

Niedenthal, P.M., Barsalou, L.W., Ric, F., & Krauth-Gruber, S. (2005). Embodiment in the acquisition and use of emotion knowledge. In L. Feldman Barrett, P.M. Niedenthal, & P. Winkielman (Eds.), Emotion and consciousness (pp. 21-50). New York: Guilford.

Oosterijk, S., Rotteveel, M., Fischer, A. H, & Hess, U. (2009). Embodied emotion concepts: How generating words about pride and disappointment influences posture. *European Journal of Social Psychology, 39*, 457-466.

Parsons, L. M. (1987). Imagined spatial transformation of one's body. *Journal of experimental psychology: General, 116(2),* 172-191.

Penn, D. C., Holyoak, K. J., & Povinelli, D. J. (2008). Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *The Behavioral and brain sciences, 31(2).*

Piaget, J. (1952). The origins of intelligence in children. New York: International Universities Press.

Proffitt, D. R., Stefanucci, J., Banton, T., & Epstein, W. (2003). The role of effort in perceiving distance. *Psychological Science, 14(2)*, 106-112.

Regier, T., & Carlson, L. A. (2001). Grounding spatial language in perception: an empirical and computational investigation. *Journal of experimental psychology: General, 130(2)*, 273-298.

Petkov, G. & Nikolova, P. (2010). Simple Pointing to Objects may Facilitate Remembering. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

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.

Rizzolatti, G., Riggio, L., Dascola, I., Umilta, C. (1987). Reorienting attention across the horizontal and vertical meridians-evidence in favor of a premotor theory of attention*. Neuropsychologia, 25*, 31–40.

Salvucci, D., & Anderson, J. (2001). Integrating analogical mapping and general problem solving: The path-mapping theory*. Cognitive Science, 25*, 67-110.

Sekiyama, K. (1982). Kinesthetic aspects of mental representations in the identification of left and right hands. *Perception & psychophysics , 32 (2)*, 89-95.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171(3972),* 701-703.

Sinai, M. J., Ooi, T. L., & He, Z. J. (1998). Terrain influences the accurate judgement of distance*. Nature, 395(6701)*, 497-500.

Spivey, M. J., & Geng, J. J. (2001). Oculomotor mechanisms activated by imagery and memory: eye movements to absent objects. *Psychological Research, 65(4*), 235-241.

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

Strack, F., Martin, L., & Stepper, S. (1988). Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology, 54*, 768–777.

Thomas, L.E., & Lleras, A. (2007). Moving eyes and moving thought: On the spatial compatibility between eye movements and cognition. *Psychonomic Bulletin & Review, 14*, 663-668.

Trudeau, J. J., & Dixon, J. A. (2007). Embodiment and abstraction: actions create relational representations. *Psychonomic bulletin & review, 14(5),* 994-1000.

Tucker, M., & Ellis R. (2004). Action priming by briefly presented objects. *Acta Psychologica, 116(2),* 185-203.

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. *Journal of experimental psycholog: Human perception and performance, 24(3)*, 830-846.

Vankov, I. (2009). Mind the gap - the cost of looking at nothing, or the performance implications of memory-induced attention shifts. In N.A. Taatgen & H. van Rijn (Eds.), Proceedings of the 31th Annual Conference of the Cognitive Science Society (pp. 1318-1323). Austin, TX: Cognitive Science Society.

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.

Wraga, M. (2003). Implicit transfer of motor strategies in mental rotation. *Brain and Cognition, 52(2),* 135-143.

Zwaan, R. A. (2009). Mental simulation in language comprehension and social cognition. *European Journal of Social Psychology , 39 (7)*, 1142-1150.

Publications

Vankov, I., & Kokinov, B. (2010) The role of action in perceiving and comparing functional relations. In S. Ohlsson & R. Catrambone (Eds.), Proceedings of the 32nd Annual Conference of the Cognitive Science Society (pp. 2081-2086). Austin, TX: Cognitive Science Society.

Vankov, I., & Kokinov, B. (2009). Grounding relations in action. In B. Kokinov, K. Holyoak & D. Getner (Eds.), New Frontiers in Analogy Research (pp. 463-472). Sofia, Bulgaria: NBU Press.

Vankov, I. (2009). Mind the gap - the cost of looking at nothing, or the performance implications of memory-induced attention shifts. In N.A. Taatgen & H. van Rijn (Eds.), Proceedings of the 31th Annual Conference of the Cognitive Science Society (pp. 1318-1323). Austin, TX: Cognitive Science Society.