NEW BULGARIAN UNIVERSITY

CENTRAL AND EAST EUROPEAN
CENTRE FOR COGNITIVE SCIENCE

Grounding relations and analogy-making in action

Ivan Vankov

A dissertation submitted in partial fulfillment of the requirements for the degree of Ph.D. in Cognitive Science

Supervisor:
Assoc. Professor Boicho Kokinov

SOFIA ● 2010

Table of Contents

Acknowledgements

I would like to thank my supervisor Boicho Kokinov for attracting me to the field of cognitive science and giving me the opportunity to work on this thesis. Boicho guided my work and supported me in every possible way. He gave me a lot of freedom and let me enjoy my graduation study from the start to the end.

I would also like to thank the members of ANALOGY team for all the discussion we had and for the useful suggestions and comments about my work.

Armina Janyan has my special gratitude for making me interested in embodied cognition and teaching me how to conduct behavioural experiments.

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 which test the predictions of the embodiment approach to relations are 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.

Chapter 1. The challenge of grounding relations

This chapter 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.

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 can not build an adequate account of relational reasoning without a psychologically plausible model of relational representations.

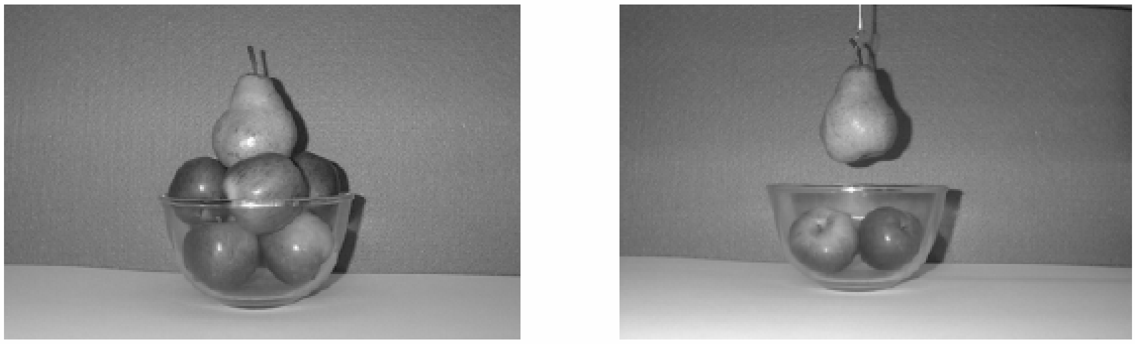

Figure 1.1. An example of the context sensitivity of relations. Most people would agree that the relation between the pear and the bowel is different in the two situations, although the positions and features of the objects are the same (example taken from Coventry & Garrod, 2004).

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 can not 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 can not 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 not define the relation between two objects without taking into account the whole situation in which the objects appear (Figure 1.1). 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.

Grounding spatial relations in perception

Analogy-making is not the only research area which is concerned with relations.

Researchers working in vision and language have also built models of relational reasoning and faced the problem of grounding relations. For example, according to the structural description theory of object recognition (e.g., Biederman, 1987), objects are represented in terms of primitive geometric components - 'geons'- organized by a set of spatial relations. Therefore models of structure-based object recognition should account for how relations between geometric components are computed by the human visual system. Vision researchers usually assume that these relations are grounded in features of the visual images. Hummel & Biederman (1992) proposed a neural network model which was able to compute simple spatial relations between geons (relative size, relative location and relative orientation) by assigning them values from corresponding scales (size, location, orientation) and then dynamically encoding the conjunctions of pairs of such values.

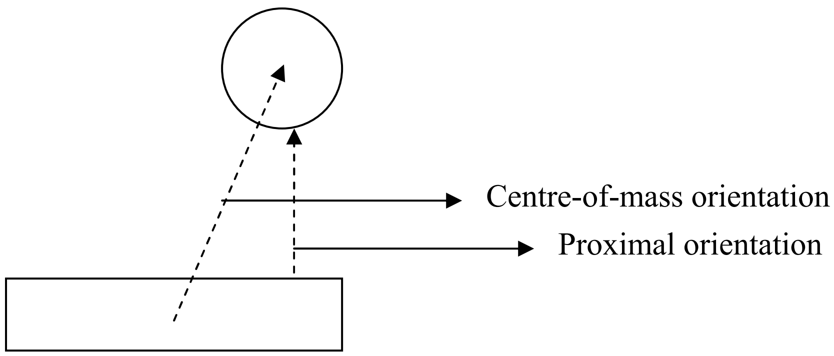

Figure 1.2. Grounding spatial language in perception. Regier (1995) proposed that the meaning of spatial terms such as 'above' is grounded in two perceptual primitives – the angles of the centre-of-mass and proximal orientations

Other researchers have been interested in how people acquire linguistic terms for relations. Regier (1995) proposed a model of grounding spatial terms in visual perception. He developed a neural network which was able to classify linguistic terms describing the spatial relation between objects by computing their centre-of-mass and proximal orientations (Figure 1.2). The model was able to explain a number of psychological findings related to the use, developmental dynamics and cross-cultural variations of spatial language. Regier, however, did not explicitly specify how the two proposed perceptual primitives are computed by the visual system. His model was also unable to account for many context-sensitivity effects, such as the one presented in Figure 1.1. Later Regier extended his model by proposing that spatial language is grounded in the process of attention and the vector sum encoding of general direction (Regier & Carlson, 2001).

The attempts to ground the meaning of spatial relations and spatial language in visual perception help to advance theories of structure-based object recognition and the origins of spatial language, but they do not provide a general solution to the problem of how relations are represented in the mind, how they are learned and how their representations affect the way we reason with them.

The role of time and action in understanding relations

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.

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 simulated cognitive agents. The simulated agents were equipped with bodies letting them interact with the world and the relational representations emerged as dynamic patterns in their behaviour while solving a relational categorization task. The parameters of the agents which managed to acquire relational knowledge were derived by evolutionary algorithms. 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 only be revealed by tracing the interplay of action and perception.

Another computational model which attempted to ground relational meaning in interactions with the environment was proposed by Möller & Schenck (2008). In their model, a simulated cognitive agent learned to predict how the visual input is changing under its movements. As a result, the agent acquired the ability to distinguish spatial configurations of obstacles (dead ends and corridors) by simulating actions and taking into account their consequences. This approach bears resemblance to previous work by Siskind (Siskind, 1994; Siskind, 2003) which proposed that the meaning of relational notions such as support, contact and attachment could be derived from perceptual input by a process of counterfactual simulation, predicting the effect of hypothetical changes to the world on the immediate future.

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.

The next chapter will present an overview of empiric research which renders support for the hypothesis that relational represenations are embodied. Chapter 3 will formulate the goals of study in this thesis. Chapters 4 and 5 will provide a computational account of how relational representations are grounded in actually executed or simulated actions.

Another implication of the embodiment view of relations is that the ability to use relations is constrained by the characteristics of the human body. We suggest that the performance in tasks involving reasoning with relations would reflect the ability of the participants to execute or simulate certain actions. We will outline several possible ways in which relational reasoning could be constrained by the characteristics of the body in Chapter 5. Chapter 6 will describe a series of experiments which aim to investigate the role of the human body in comparing relations.

Finally, the embodiment of relations has implications for the development of relational knowledge. An embodiment view of relations entails that the formation of relational concepts and the learning to reason with relations is dependent on the acquisition of specific motor skills. We will not address the issue of learning relations in this thesis, but we will return to it in the final Chapter 7 where the limitations of the proposed account of relational thinking and directions for future studies are discussed.

Chapter 2. The embodiment view of relations

The goal of this chapter is to introduce 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.

Theories of embodied cognition

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. According to Barsalou, concepts are represented by 'simulators' – association areas in the brain which are able to reactivate patterns of perceptual and motor information. A simulator can produce limitless 'simulations' which will represent instances of the concept. In this account, categorization is viewed as the ability to simulate a perceived entity. Such a view is consistent with our idea, presented in the first chapter, that the recognition of relations involves the execution or simulation of actions. Barsalou did not provide any examples of how relational concepts and relational reasoning could be accommodated in his theory, though he argued that all conceptual knowledge is grounded in simulations.

The PSS theory predicted a number of effects for which empiric support was found. For example, one of the predictions is that conceptual represenations are essentially modal. Pecher, Zeelenberg & Barsalou (2003) reasoned that if this is true than there must be a processing costs of switching between modalities when verifying the properties of concepts. They found out that people perform faster when they have to verify a property in the auditory modality (e.g., BLENDER-loud) after verifying a property in the same modality (e.g., LEAVES-rustling) than after verifying a property in a different modality (e.g., CRANBERRIES-tart). In another study, Richter and Zwaan (2009) showed that participants respond faster in a lexical decision task when preceding colours matched the colour named by the word and colour-discrimination responses were slowed down when preceding colour words mismatched the test colour. The predictions of the PSS theory were also supported by neuro-imaging studies. For example, Simmons, Martin & Barsalou (2005) found that pictures of appetizing foods activate gustatory cortices for taste and reward.

The idea that cognition is grounded in simulation has also been exploited in language research (for reviews, see Fischer & Zwaan, 2008; Zwaan, 2009). Zwaan (2003) proposed that reading or hearing words activates experiential representations of their referents, including perceptual, motor and emotional information. In this view, the linguistic description of a situation results in mental simulations similar to the result of directly experiencing the situation. A significant body of evidence was accumulated in support of the hypothesis that comprehending language involves running perceptual-motor simulations (e.g., Stanfield & Zwaan, 2001; Glenberg & Kaschak, 2002; Richardson,. Spivey, Barsalou & McRae, 2003; Kaschak et al, 2005, Zwaan & Taylor, 2006; Kaup et al. 2007, Yaxley & Zwaan, 2007). For example, Stanfield & Zwaan (2001) demonstrated that people mentally represent the orientation of an object implied by a verbal description. Participants were asked to read a sentence that implicitly suggested a particular orientation for an object (e.g. 'He hammered the nail in the wall' implies a horizontal orientation of the nail) and then determine whether a picture referred to any of the objects mentioned in the sentence. Recognition latencies were shorter when the orientation of the pictured object matched the implied orientation of the object mentioned in the sentence (for example, a picture of horizontally oriented nail was recognized faster after listening to the sentence 'He hammered the nail into the wall' than after the sentence 'He hammered the nail into the floor'). The results were interpreted as evidence that people run sensorimotor simulations when comprehending the verbal description of the scene. In a similar study, Glenberg and Kaschak (2002) asked participants to judge whether sentences were sensible by making a response that required moving toward or away from their bodies. When a sentence implied action in one direction (e.g., "Close the drawer", implies action away from the body), the participants had difficulty making a sensibility judgment requiring a response in the opposite direction. Evidence that language understanding involves perceptual-motor involves simulations has been also found in eye-tracking (e.g., Richardson & Matlock, 2007, Coventry et al, 2009) and neuro-imaging studies (e.g., Pulvermüller, Haerle, & Hummel, 2001; Willems, Hagoort, & Casasanto, 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 & 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).

Neural mechanisms underlying perceptual-motor simulations

The discovery of the mirror neuron system in the ventral premotor cortex (area F5) of macaque monkeys (Rizzolatti & Cragehero, 2004) shed light on the possible neural mechanisms underlying the perceptual-motor simulations. A mirror neuron is a neuron that fires both when an animal acts and when the animal observes the same action performed by another. It is believed that mirror neurons underlie action understanding and the ability to understand and imitate other's behaviour by simulating it internally (Hurley, 2008). An interesting finding, which accompanied the discovery of mirror neurons, was that there is another class of normal (canonical) neurons – the 'visuomotor' neurons – which were located in the same area of the brain of the monkey and fired during the execution of grasping actions

(Rizzolatt, Fogassi, Gallese, 2001). However, visuomotor neurons that selectively discharged during grasping of an object also discharged selectively to the mere presentation of that object, when the monkey was not expected to do anything. The majority of visuomotor neurons showed selectivity for one or few objects. This finding renders support for the hypothesis that object representations include motor programs and the recognition of objects may involve running sensorimotor simulations.

The existence of mirror neurons in humans is still an open question (Turella, Pierno, Tubaldi & Castiello, 2009; Lingnau, Gesierich & Caramazza, 2009), but several studies have provided indirect evidence for it (e.g., Buccino et al, 2001; Pulvermüller, Haerle, & Hummel, 2001). Participants in the Buccino et al (2001) fMRI study were presented with videotaped object- and non-object-related actions performed with different effectors (mouth, arm, or foot). Somatotopically organized activation was found in the premotor cortex, thus providing support for the existence of a mirror system in humans. Even if it finally turns out that such a system exists, there is no conclusive evidence hat it is involved in action understanding and subserves embodied object representations (Negri et al., 2007; Hickok, 2009).

Another line of research aiming to elucidate the neural mechanisms underlying embodied cognition is the dissociation of dorsal and the ventral visual processing streams (Goodale & Milner, 1992). The dorsal visual processing stream subserves object-directed action, whereas the ventral visual processing stream subserves visual object recognition. It has been shown however that the dorsal system affect performance in a visual object recognition task, but only for a category of tool-like objects (Almeida, Mahon, Nakayama & Caramazza 2008). This finding implies that the recognition of tools depends on the activation of motor programs related to their use through the dorsal visual stream. The role of dorsal system could be to quickly compute motor-relevant information before the object is actually recognized and enable subsequent recognition of the object by means of perceptual-motor simulations

(Almeida, Mahon & Caramazza, in press).

Evidence for embodiment of relations

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. The goal of the second part of this chapter is to review existing evidence for the embodiment of some kinds of relations.

Spatial terms and spatial relations

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 (Regier & Carlson, 2001). Logan (1994) conducted a series of experiments in which he asked participants to search for a target in a field of distractors which differed from the target only in the spatial relation between their parts. The results indicated that the discovery of spatial relations requires attention. In a later study Logan (1995) found that linguistic cues such as 'above', 'below', 'left' and 'right' are used to direct attention from one object to another. Regier and Carlson interpreted these results as evidence that spatial relations are grounded in attention in the sense that the encoding of spatial relations is dependent on the direction and magnitude of attention shifts. 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.

Recent eye-tracking studies have rendered support for the hypothesis that spatial relations are grounded in eye and head movements. Demarais & Cohen (1998) asked participants to solve transitive inference problems with the relational terms left/right and above/below while their eye movements were recorded by electrooculography. The results revealed that people made more horizontal and fewer vertical saccades while solving problems with the left/right terms than while solving identical problems with above/below.

In another study, Spivey and Geng (2001) recorded the eye-movements of participants while they faced an empty computer screen and listened to pre-recorded descriptions of spatiotemporally dynamic scenes. The experiment revealed that people tended to make saccades in the same direction as the spatiotemporal dynamics of the auditorily presented scene descriptions. For example, they tended to move their eyes upwards when they listened to a story asking them to imagine that they were standing across the street from tall building and describing a series of events happing on subsequent floors of the building, starting from ground floor to the top one.

Coventry and et (2009) recorded eye-movements while subjects made judgments of how well linguistics descriptions, including the spatial prepositions 'above'/'over' and 'below'/'under', described the spatial relation between two objects with varying functional relations. The results indicated that the use of spatial language differentially directs attention during examination of a visual scene. Also, it was found that linguistic variations are correlated with specific eye-movement patterns.

The results of the eye-tracking studies imply that the representations of spatial relations involve specific patterns of eye-movements. It remains unclear however what exactly the role of eye-movements is – whether they are necessary for thinking about spatial relations or they are just a side effect. It is difficult to answer this question as long as there is no efficient way of manipulating eye-movements, neither there exist such lesions that selectively impair eye-movements.

We will present a computation model of grounding spatial relations in Chapter 5, according to which eye-movements are essential for encoding and reasoning with spatial relations.

Distance and slant perception

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 idea that the perception of distance is embodied dates back to the work of the philosopher George Berkely (1709). He noted the that projection of a point of light into the eye conveys no information about egocentric distance to an object and suggested that it is perceived in terms of the effort required to walk the way to the object. Proffitt, Stefanucci, Banton & Epstein (2003) empirically tested this hypothesis by asking participants to make metric distance judgements either unencumbered or wearing or while wearing a heavy backpack. The results were that the distance judgments were greater for the latter group than for the former. A similar result was obtained when the manipulation was a visual-motor adaptation that reduced the anticipated optic flow coinciding with walking effort.

In another study of distance perception, Sina, Ooi & He (1999) found that the distance judgements are influenced by terrain properties. People jugged the distance to remote objects as bigger when there were gaps and obstacles between themselves and the objects. This implies that the participants in this study simulated walking to the objects in order to assess how far they are.

Linkenauger et al. (2009) found that another factor which affects the perception of egocentric distance is the reachability of objects. Participants estimated the distances to tools with handle orientations that made them either easy or difficult to grasp with their dominant and non-dominant hands. Right-handed participants perceived tools that were more difficult to grasp to be farther away than tools that were easier to grasp. No such difference was found for left-handed participants. The authors concluded that 'the perception of spatial layout is action specific and is scaled by the body's abilities to perform intended actions'. The fact that a difference was found between left-handed and right-handed subjects different supports the hypothesis that cognition, and in particular the encoding of metric spatial relations, is constrained by the characteristics of the human body.

Studies of geographical slant perception found similar effects as in distance perception. Bhalla & Proffitt (1999) showed that hills appear steeper to people who are encumbered by wearing a heavy backpack, are fatigued, are of low physical condition, or are elderly and/or in declining health.

The above presented studies imply that people simulate actions when trying to compute the metric properties of spatial relations. These findings are in support of our proposal that that relations and relational reasoning are grounded in physically executed or simulated interactions with the environment.

Mental imagery

Mental imagery is defined as "an experience which resembles perceptual experience, but which occurs in the absence of the approatiate external stimuli" (Stanford Encyclopaedia of Philosophy, online resource). It has been argued that mental imagery utilizes the same neural circuits which are used by the perceptual and motor systems (Kosslyn, 1994) and therefore mental images can be considered as a kind of perceptual-motor simulations in the sense Barsalou proposed (Barsalou, 1999).

The connection between mental imagery and relational knowledge is in that some of the most poplar 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. The same effect was found for two dimensional shapes as well (Cooper, 1975). These findings were interpreted by assuming that people mentally rotate one of the objects until its visual image matches the image of the other object.

In another study, Bundesen & Larsen (1975) asked participants to compare the shapes of different-sized shapes. Again it was found that the reaction time for correct responses to test pairs of figures of the same shape and orientation increased approximately linearly as a function of the linear size ratio of the figures. The authors concluded the task was performed by a gradual process of mental size transformation of one of the members of each pair of figures to the format of the other one.

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). For example, Sekiyama (1982) found that performance in such tasks reflect specific anatomical constraints of the joints involved in hand movements. The findings of Sekiyama and Parsons were also supported by brain-imaging studies (e.g., Wraga, 2003; Hanakawa, Hosoda, Shindo & Honda, 2007).

More recent studies have found evidence that the transformations of visual images are guided by motor processes even in the case of images of abstract objects rather than of body parts. Wexler (1999) asked participants to solve the classical two-dimensional mental rotation task, while executing an unseen motor rotation in a given direction. He found that performance in the mental rotation task was better when the direction of mental rotation was compatible with the direction of motor rotation. The fact that executing a real rotation interferes with mental rotation implies that mental transformations are actuated by the motor systems of the body. In another study, Flusberg, Jenkins & Boroditsky (2009) asked the question whether objects that are more difficult to physically manipulate are also more difficult to mentally manipulate. Participants interacted with two wooden objects modelled after the figures from Shepard and Metzler's (1971) classic mental rotation study. One object was easy to physically rotate while the other was difficult to rotate. They then completed a mirror-image mental rotation task consisting of images of the manipulated objects. Participants were slower to solve the mental rotation task for trials consisting of images of the hard-to-rotate object. The result however was obtained only when participants were explicitly asked to use a motor strategy in the task.

The studies of mental rotation and mental size transformation suggest that the meaning of the 'sameness' relation is grounded in simulated transformations of the visual input. In Chapter 6 we will present an experiment which addresses a specific prediction of this assumption.

Higher-order relations

Higher-order relations express relations between relations. They are believed to lie at the core of abstract problem solving, including analogy-making (Gentner, 1983). 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.

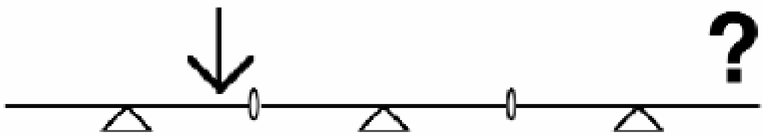

Figure 2.1. An example of a connected balance beam problem (Dixon & Dohn, 2003). Three balance beams are connected by flexible joints (shown as ovals). The arrow indicates that the right arm of the leftmost balance beam will be pushed down. The participants were asked to predict the motion of the right arm of the rightmost balance beam.

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. Dixon & Dohn (2003) asked participants to solve a set of unfamiliar problems about a series of connected balance beams. The beams were connected by flexible joints and were shown in a linear series (Figure 2.1). For each problem, the participant was told that the right arm of the first beam in the series would be pushed either up or down. The participants were asked to predict the resulting movement of the final beam in the series (whether it would go up or down. Half of the participants were asked to solve the problems but were not instructed on a strategy. On the basis of pilot work, it was known that most of these participants would solve the problem by physically simulate the movement of the beams. The other half of the participants was given a formal rule for solving the problem. In the second part of the experiment, the participants had to solve another task which was analogous to the first one in that both of them required the discovery of the alternation relation. The results indicated that the participants who had to discover the alternation relation by simulating the movements of the beams in the first problem transferred the relation to the second problem more quickly and used it more consistently than did participants who had been given explicit instruction about the relation.

In a subsequent study (Trudeau & Dixon, 2007), the same data were reanalyzed and it was discovered that the number of alternating actions in episodic memory prior to discovery of the alternation relation predicted its generalization to new problem types. The authors concluded that these results show that actions can function as the representational substrate of higher-order relations.

Implicit transfer of motor strategies

Day & Goldstone (2009) reported an experiment which found similar effects as Dixon & Dohn (2003). Day and Goldstone first asked the participants to interact with a simulation of a physical system involving the oscillating motion of a ball that was suspended between two elastic bands. Next, all individuals participated in another task, which was much more conceptual and involved regulating the population of a city. While this second simulation differed considerably from the first, both in terms of its content and its visual display, the system was governed by the same underlying principles as the first task. In both tasks, participants were asked to accomplish a specific goal, which required the development of an appropriate strategy. The manipulation was in the relationship between the goals for the two tasks. For some participants, the two goals were analogous, and thus required analogous strategies to achieve. For other participants, the two goals were structurally dissimilar, with each requiring a unique strategy. The main result was that participants required reliably fewer trials to complete the second, conceptual, task when it required achieving a goal that was analogous to that of the physical system simulation task. This finding suggests that discovering relational similarity might be facilitated by procedural knowledge and serves as indirect evidence that relational knowledge is grounded in action and perception.

Embodied analogical mapping

The previous paragraphs described experiments which suggest that analogical transfer is grounded in procedural knowledge. There is also evidence that another sub-process of analogy-making – analogical mapping – is embodied (Salvucci & Anderson, 2001; Gordon & Moser, 2007). Gordon & Moser (2007) recorded eye-movements of participants while solving picture analogies in which they had to identify the object in one picture that "went with" an object in another, simultaneously presented picture. The results showed that the pattern of saccades between objects was a very sensitive measure of the time course of both relational and object-matching processes.

Detecting perceptual-motor similarities between analogous situations

An embodied account of relations and relational reasoning would predict that comparing two relations would be faster if the relations are grounded in overlapping perceptual and motor patterns of activation. Clement (2004; 2009) reports evidence from expert reasoning protocols that when evaluating analogies, people use to simultaneously run imagistic simulations of the two analogous cases in order to detect perceptual and motor similarities between them. We shall return to this problem in the next chapters and will argue that the ability to compare the dynamics of perceptual-motor simulations is fundamental to the process of relation comparison.

Conclusion

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 current thesis are twofold. First, it has to propose and describe a computational model of grounding relations and relational reasoning in the sensorimotor dynamics resulting from physically executed or simulated interactions with the environment. The model should be able to account for existing empirical findings, such as the ones discussed in the previous chapter. 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 description of the computational model is given in chapters 4 & 5. It consists of a brief overview of the modelling environment and two simulation studies which demonstrate the operation of the model and materialize its predictions.

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 6 describes a series of experiments which aim to fulfil these goals.

The results of the computational and the experimental studies will be unified together in the last chapter of this thesis. We will discuss to what extent the presented computational and empirical work solves the problems of grounding relations. We will also discuss the problems and limitations of the proposed approach and will outline directions for future studies.

Chapter 4. AMBR

This chapter 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 in tact with the general principles and philosophy of AMBR.

Why AMBR

There are several reasons why the AMBR was selected. 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.

Basic Principles

AMBR (Associative Memory Based Reasoning) 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' (Figure 4.1). '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 (Figure 4.2). 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.

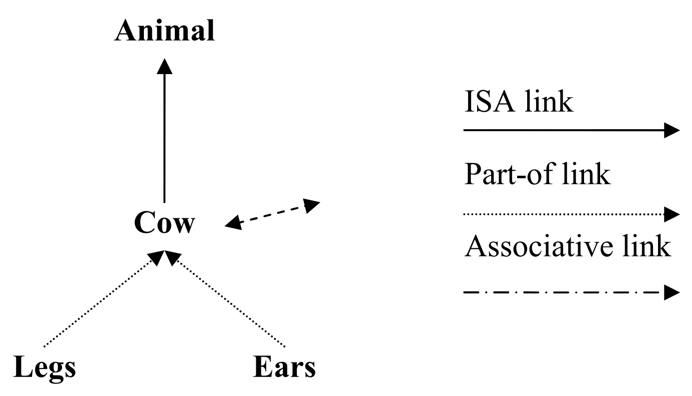

Figure 4.1. Knowledge representation in AMBR. There are three types of links between nodes.

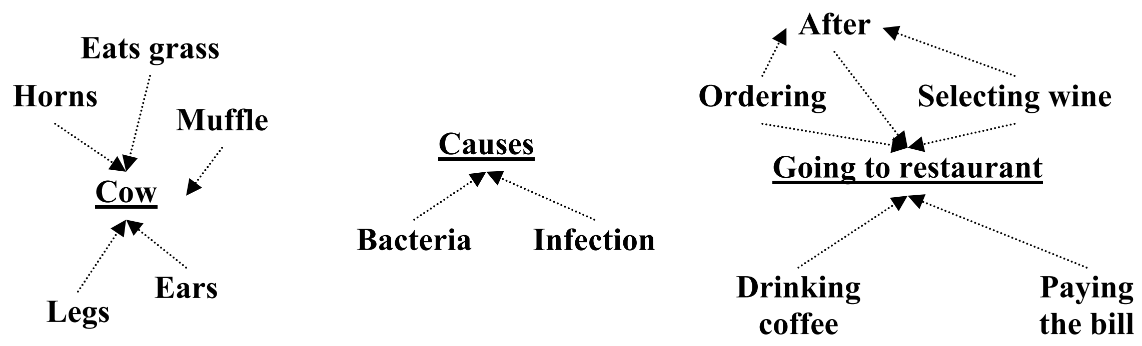

Figure 4.2. Exampels of bindings nodes: a concept (to the left), an instance of a relation (middle) and script of an event (right).

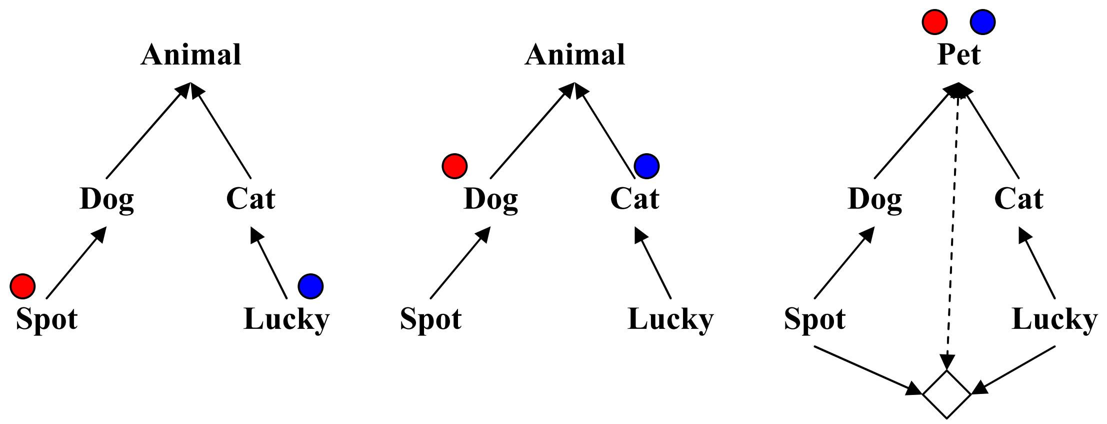

Figure 4.3. Establishing semantic similarity by marker passing. Instances of concepts send markers along the 'is-a' links. When two markers intersect at a a node, a hypothesis node is created. Its represents the mapping of nodes. The activation of the hypothesis node is a function of the activation of the nodes that were mapped, as well as the node at which their markers have intersected.

A hypothesis node is created when a semantic similarity is established. Semantic similarity is dynamically computed by a process of marker-passing (Figure 4.3). Hypothesis nodes can also be created due to top-down or bottom-up structural constraints (Holyoak & Thagard, 1989). The hypothesis nodes are created by using local information only – there is no central mechanism which monitors the whole network of nodes in AMBR. Hypothesis nodes which represent consistent mappings support each other by positively weighted associative links and inconsistent hypothesis inhibit each other (Figure 4.4). Thus the hypothesis nodes create a constraint satisfaction network. The outcome of cognitive processes such as memory retrieval, recognition, categorization and analogy-making is determined by the resolution of the constraint satisfaction network

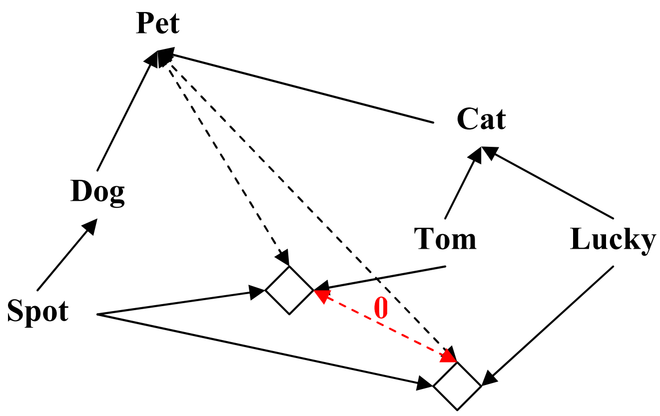

Figure 4.4. Constraint satisfaction network - ihhibition of rivalling hypothesis nodes.

The role of binding nodes is to bind together elements of a distributed representation (Figure 4.2). Binding nodes are created in two ways – bottom-up and top-down. Bottom-up creation of binding nodes assumes that the perceptual input is already organized by automatic processes which lie outside of the scope of the AMBR model. Top-down creation of binding nodes is implemented by a process of analogical transfer (Figure 4.5). Competing sources of analogical transfer can create binding nodes which are inconsistent. Inconsistent binding nodes start to inhibit each other and also form constraint satisfaction networks. In other words, there could be alternative interpretations of the perceptual input which compete with each other. We have used these mechanisms to model the recognition of ambiguous pictures and generate a non-trivial prediction, for which empirical support was found (Kokinov, Vankov & Bliznashki, 2009).

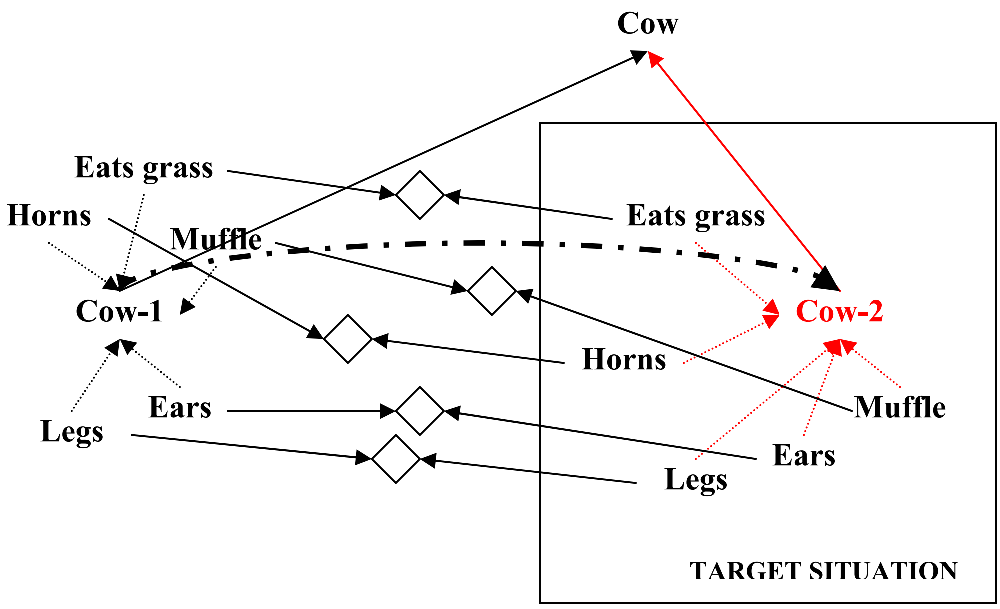

Figure 4.5. Analogical transfer of binding nodes. The transferred elements are drawn in red. The thick dotted line depicts the transfer process and is not part of the representation. Elements are transffered only to the target situation.

The input to an AMBR simulation consists of a number of nodes which are connected to a source of permanent activation. The input nodes are also marked as targets so that the system tries to map them to other nodes retrieved from memory. The set of input and target nodes can be changed in the course of the simulation.

Representation of actions

The most important extension of AMBR is the introduction of 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 (Figure 4.6). The idea is that the activation of a motor program can not 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.

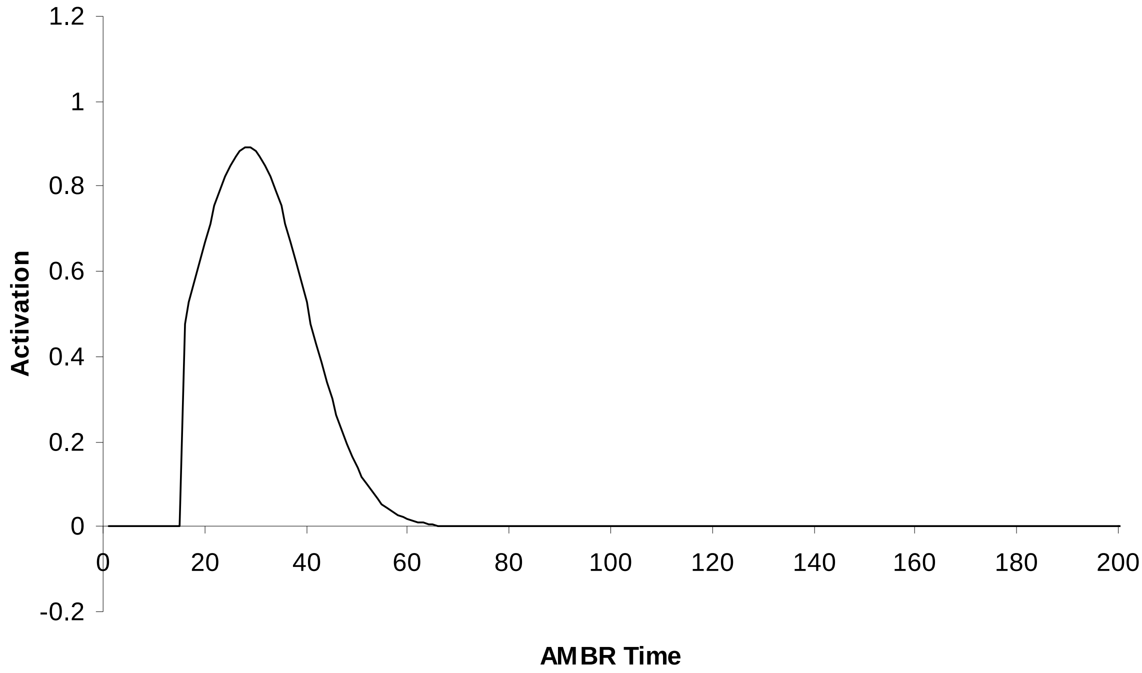

Figure 4.6. The activation profile of an action anode created at time 15 (in AMBR time units). The activation of the action node grows to its maximum and then grdually decreses, irrespective of how much activation it receives from other nodes.

The next chapter described how the extended version of AMBR presented in this chapter is used to model the grounding of relations and relational reasoning in action.

Chapter 5. Simulations

The goal of this chapter is to provide a more detailed description of the idea that relations are grounded in action and to formulate the predictions it generates. To this end, we conducted two simulations based on the AMBR model of analogy-making described in the previous chapter.

The goals of both simulations are to propose psychologically plausible mechanisms of embodied relational reasoning and to generate testable predictions. The simulations are deliberately simplified in order to focus on the main contributions of the model.

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.

Simulation 1

The goal of Simulation 1 is twofold. First, it is designed to serve as a general demonstration of the role of action in perceiving and comparing relations. Second, it aims to demonstrate how the embodied view of relations could solve one of the pivotal problems of relational reasoning – the role-filler binding problem (Hummel, 1999).

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.

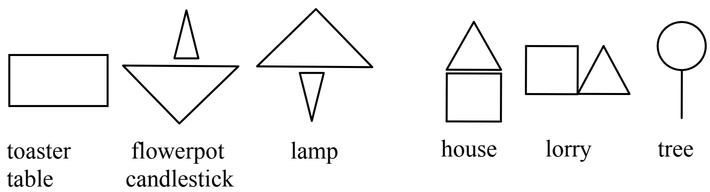

Figure 5.1.1. RecMap knowledge domain (Petkov & Shahbazyan, 2007)

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 5.1.1). 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.

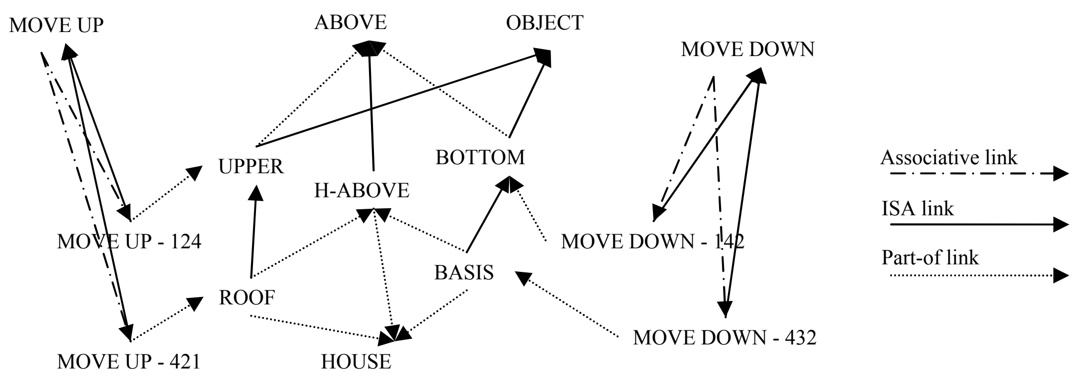

Figure 5.1.2. Representation of a spatial relation (ABOVE) and an instance of it (H-ABOVE).

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 representation of 'above' is schematically depicted in Figure 5.1.2.

Whenever a certain action is executed, the attention focus is moved in the corresponding direction. For example, when the hypothesized cognitive agent '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’.

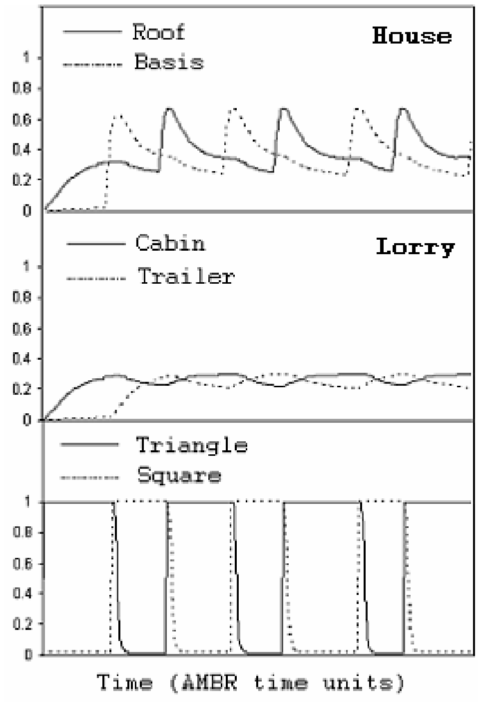

Figure 5.1.3. The dynamics of the sensory-motor loop. The perceptual input is displayed in the lowest part. Note the oscillating activation levels of the two house parts (top panel) which coincide in time with the changes of the perceptual input (bottom panel). The activation dynamics of the lorry parts (middle panel) are not synchronized with the perceptual input.

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.

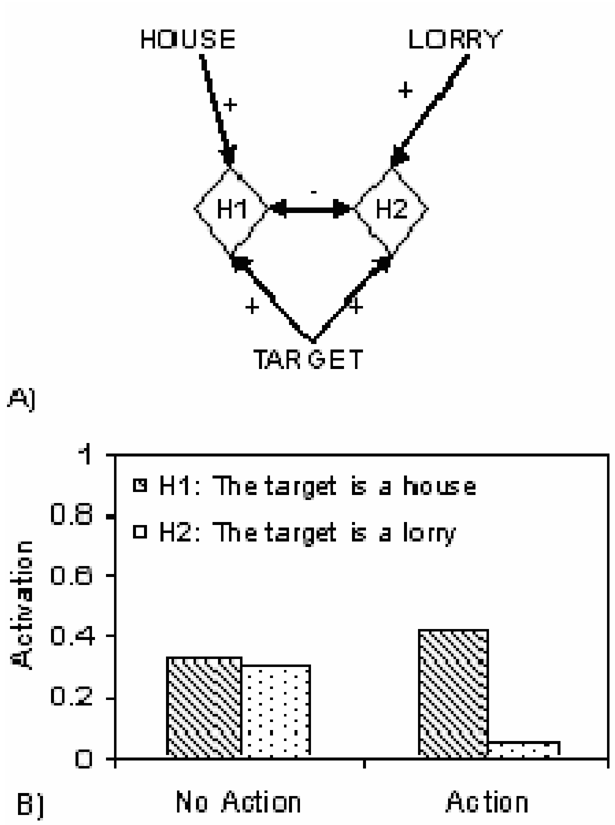

Figure 5.1.4. a) The rival recognition hypotheses characterizing the target scene. Note he inhibitory link between them. b) The activation levels of the recognition hypotheses after several movements are performed and when no action is modeled, but just the two target objects are changing turns as attentional focus. The maximum activation level is 1.0.

An action is executed by creating a new action node (f.e. MOVE-UP-384) which points to the corresponding motor concept node (MOVE-UP). The activation of MOVE-UP-384 is not persistent, but it retains some activation for a period time according to the decay rate of its activation function. The activation of MOVE-UP-384 spreads throughout the network to other instances of MOVE-UP, such as MOVE-UP-124, MOVE-UP-421 and so forth (Figure 5.1.2). Thus, 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 (f.e. UPPER) as well as other objects in memory that had been bound to the same relational role in the past (f.e. ROOF). 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 5.1.3).

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. Figure 5.1.4 shows the outcome of the simulation with embodied relations. It is clear that when no actions were performed, the system was able to disambiguate the perceptual input. It is notable that just a few movements were enough to resolve the constraint satisfaction network – a fact that is consistent with our intuition that we do not need to wonder much in order to establish a certain spatial relation.

Discussion

Simulation 1 demonstrated how an embodied approach to representing relations would work. Although the simulation was based on toy examples, its dynamics was 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’.

Simulation 1 also proposed a solution to the role-filler binding problem. Initially the candidate arguments of a relation were mapped to both relational roles but in the course of the simulation a pair of bindings become dominating because it was consistent with the dynamics of the action execution. Thus the relational role of an object is determined not by any specific features it may inherently posses, but by its role in the executed or simulated actions which underlie the meaning of the relation.

The presented approach to modelling representing and reasoning with spatial relations predicts that the comprehension of spatial relations is dependent on the motor activity resulting from executing actions related to attentional control. This prediction is supported by the results of sensory substitution studies (e.g. Collins & Bach-y-Rita, 1973), which have shown that attentional control is essential for space perception, even more important than perceptual resolution. There is also evidence that the comprehension perception of spatial relations and spatial terms recruits attentional resources (Logan, 1994; 1995). On the other hand, eye-tracking studies (e.g. Spivey & Geng, 2001) have shown that people tend to move their eyes when imagining spatial configurations. All this results renders support to our hypothesis that actions related to attentional control, such as eye and head movements, are employed in the representation of spatial relations. There are also other models of spatial relations which generate the same prediction. For example, according to the Attention Vector Sum (AVS) model (Regier & Carlson, 2001), spatial relations 'are grounded in the process of attention and in vector-sum coding of overall direction'. The AVS model successfully accounts for a number of linguistic phenomena related to the use of spatial terms. The proposed approach to grounding spatial relations is completely compatible with the AVS model.

Simulation 2

Simulation 1 demonstrated how relations could be grounded in actually (physically) executed actions. The encoding of spatial relations was entirely bottom-up driven. We assume that this is the case for all relations as ubiquitous as spatial relations, for which the actual execution of grounding actions comes at no cost. There are however cases in which the actual execution of the actions underlying relation meanings is not possible. We suggest that in these cases the execution of relevant actions is simulated.

The goal of Simulation 2 is to show how relations are grounded in simulated interactions with the environment. 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 idea of affordances was 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. Another way to think about affordances is to regard them as 'invitations for action'. A number of studies have shown that the mere perception of an object immediately activates potential motor interactions with it (Beauchamp et al., 2002; Beauchamp Martin, 2007; Adamo Ferber, 2008; Buccino et al., 2009). It is believed that this activation of motor information is automatic and subserves object recognition. In our view perceived affordances also contribute to discovering relations between objects. We will refer to such affordances as relation affordances.

Figure 5.2.1 depicts the representation of a relation affordance in AMBR. There are two objects with a bundle of attributes and relations which all activate the execution of an action. The action is represented by the binding node of the whole representation (all other nodes have part-of links to it). If there are two objects in the target situation, a lock and a key, and the key is in front of the lock, then they are mapped to the corresponding nodes of the relation affordance and the action node is transferred to the target situation (Figure 5.2.2).

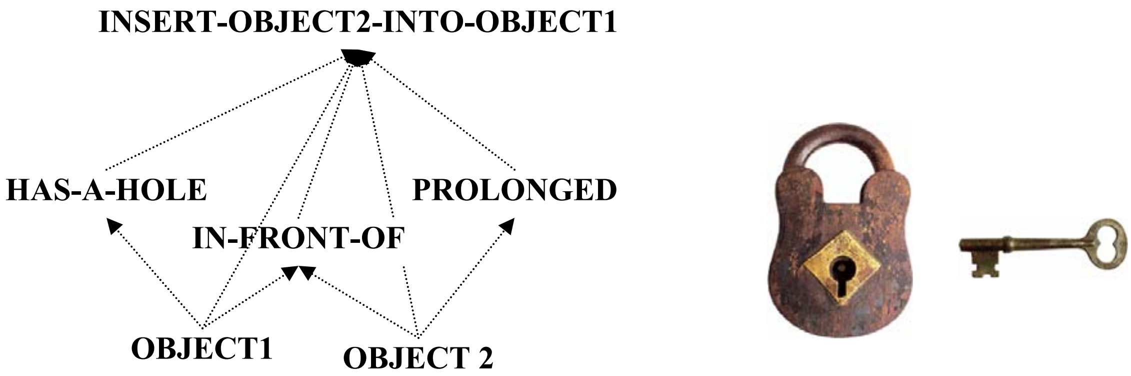

Figure 5.2.1 An example of the representation of a relation affordance. There are two objects which are spatially arranged in such a way that they invite the execution of a specific action. For example, a key in front of a lock invites inserting the key into keyhole.

The final result of the recognition of a relation affordance is the appearance of a transferred action node in the target situation. This action node represents the activation of (potentially irrelevant) motor information. Transferring the action node does not necessary entail that corresponding node would be executed. If its activation is not high enough, or if inhibitory mechanisms are in play, then the action would not be actually executed. However, the new action node can participate in other operations and lead to transferring other structures in the target situation as long as its activation exceeds the working memory threshold.

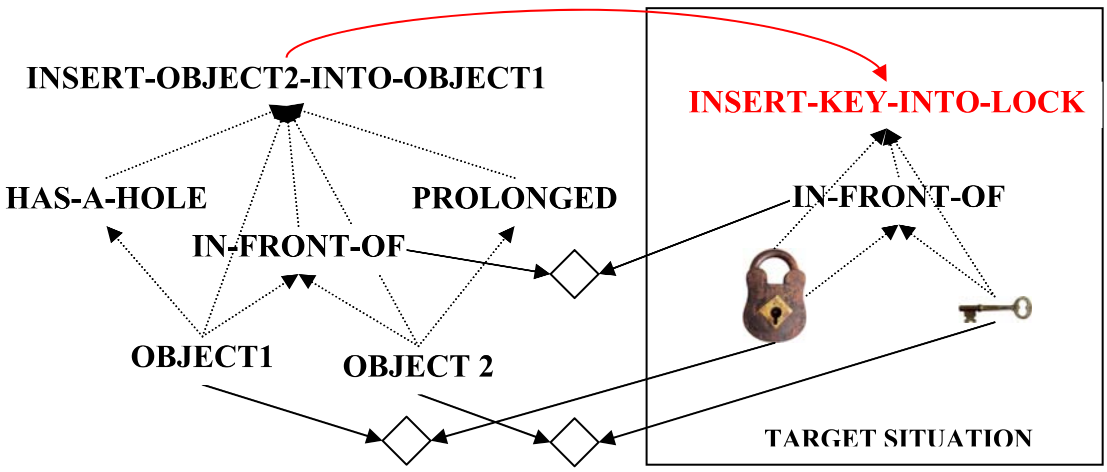

Figure 5.2.2 The processing of an affordance. The objects and spatial relations comprising the target situation are mapped to the representation of the affordance. As a result, the relevant action node is transferred in the target situation. Transferred elements are coloured in red.

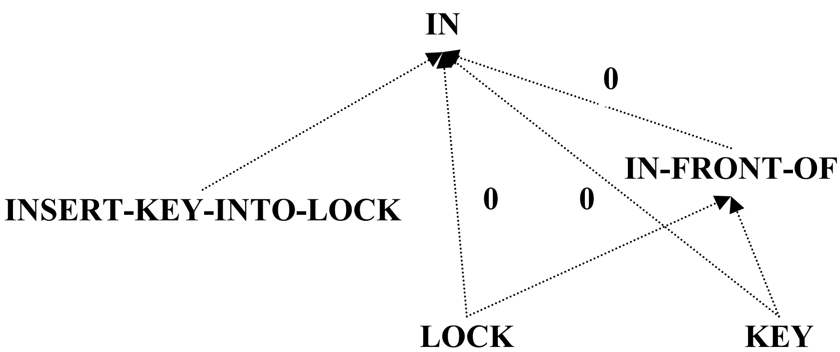

Figure 5.2.3 Representation of transformational knowledge. The diagram represents knowledge about what the spatial relation between a key and a lock will be if a key is inserted in the keyhole of the lock. INSERT-INTO is an action node. The top binding node is an instance of a spatial relation. Note the zero weights of all the links pointing to the binding node except for the link from the action node.

The second notion that we need to introduce in our modelling environment is transformational knowledge. Transformational knowledge is knowledge about how the perceptual input is changed by action execution. For example, transformational knowledge can predict how a lock and a key will look like if we insert the key into the keyhole of the lock. Transformational knowledge is a kind of sensorimotor knowledge, which supposedly links together elements of our perceptual and motor experience.

Figure 5.2.3 shows how transformational knowledge is represented in AMBR. The idea is that the execution of an action can create new object attributes and relations between objects. The spatial relation which is at the top (the top binding node) is transferred to the target situation only when all the others nodes are already mapped (Figure 5.2.4).

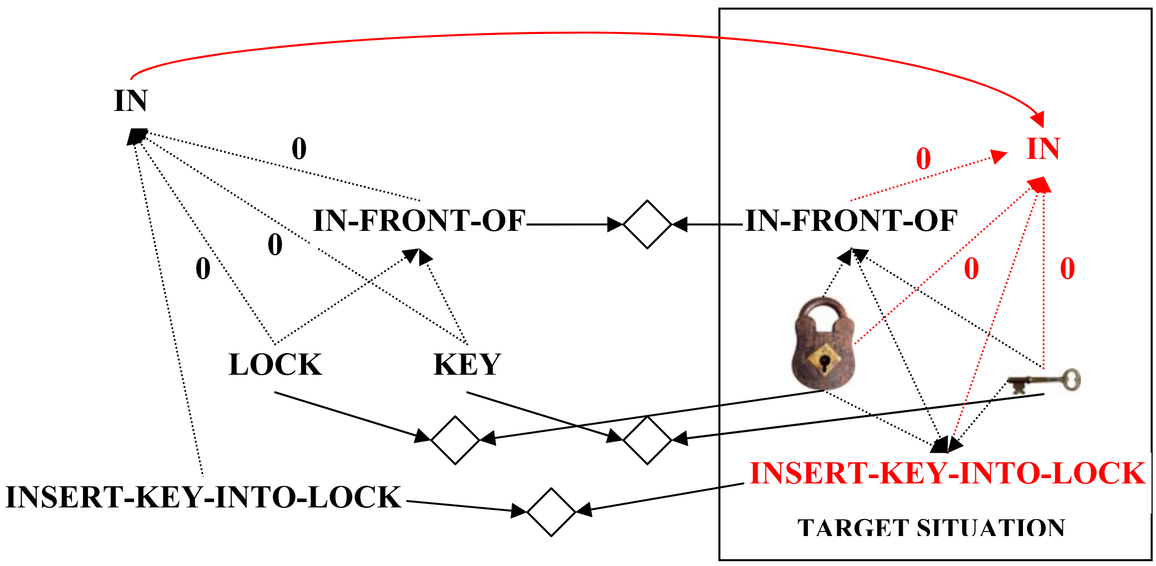

Figure 5.2.4 Application of transformational knowledge. The simulation of an action execution in a given situation leads to transferring alternative descriptions of the target scene. The transferred action nodes (red INSERT-INTO in this example) are the only source of activation for the new structures (red IN).

All the part-of links pointing to the top binding node have zero weights, except for the link from the action node. Such a configuration ensures that the activation of the relation which resulted from the action execution simulation is entirely dependent on the activation of the action. When the activation of the simulated action fades away, the activation of all the structures that were created as a result of it also fade away. Thus, the effect of an action which is only simulated, but not actually executed, is temporary.

The presented approach to modelling transformational knowledge has two key aspects. First, the simulation of an action leads to encoding new relations in the target situation. In this way, it is possible to foresee the perceptual consequences of certain actions without actually executing them. The second key point is that the effect of the simulated actions is temporary.

Encoding and comparing relations by running sensorimotor simulations

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. The two mechanism described above can be employed to model such dynamic recognition and comparison of relations.

The recognition of a 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.

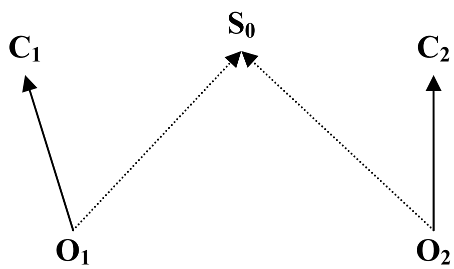

Figure 5.2.5. The input to Simulation 2. The target scene consisted of two objects O_1 and O_2, instances of the concepts C_1 and C_2. S_0 described the initial state of the target scene.

For example, in order to comprehend the functional relation unlock(key, lock), we go through simulating how we grasp the key, move it to the lock, insert it into the keyhole, turn the key clockwise. When we reach this point the transformational knowledge predicts that in the next stage the lock would be unlocked. At this point we realize that the key could unlock the lock and we discover the functional relation between these two objects.

We conducted a computational study in order to demonstrate the dynamics of relational categorization. The input to the simulation was a scene consisting of two objects O_1 and O_2. A single binding node S_0 described the initial state of the scene, i.e. the spatial relations between the objects and all other information which contributed to their visual appearance. We decided to encode the initial state of the scene by a single notation in order to simplify the simulation and ease its exposition.

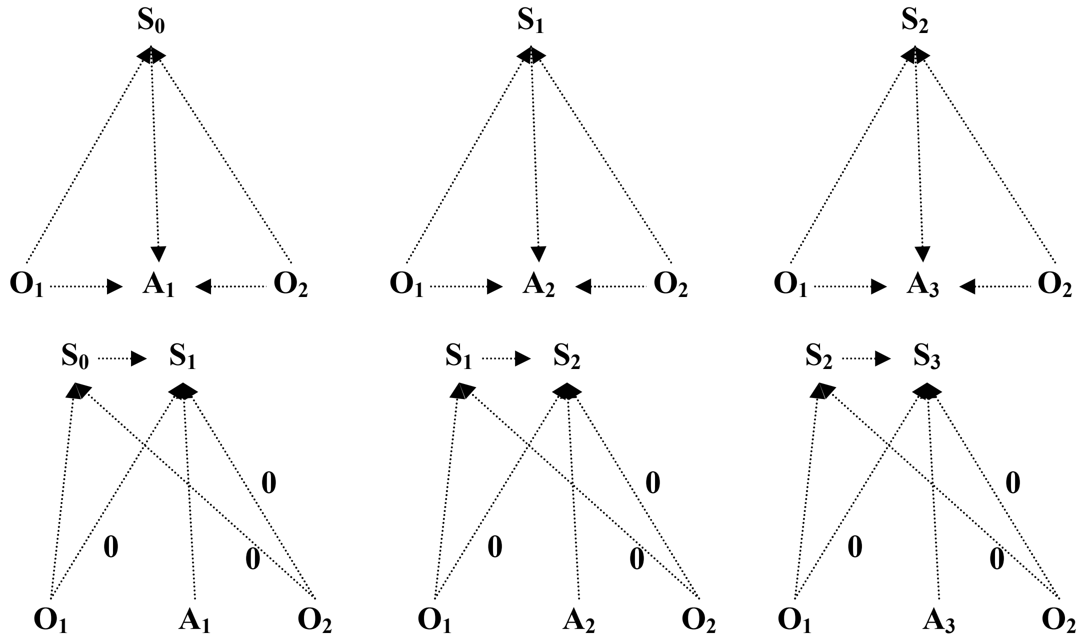

Figure 5.2.6. The contents of AMBR memory in Simulation 2 (top row- affordances, bottom row – transformational knowledge).

The definitions of affordances and the transformational knowledge were modelled as AMBR memory traces (Figure 5.2.6). The definition of the relation itself is given in Figure 5.2.7. It consists of the initial state S_0 and the goal state S_3 which resulted from applying affordance-drive transformation to the initial scene.

Figure 5.2.8 presents the dynamics of the cognitive processes which were modelled. The recognition of affordances led to simulating the executions of actions (A_1, A_2, A_3). The simulated actions temporary reorganized the perceptual input by encoding new relations and attributes (summarized by state nodes S_1, S_2, S_3). The relation between O_1 and O_2 was recognized when the simulated transformations of the target scene reached the goal state S_3. The representation of the functional relation remained active when the effect of the simulated actions had faded away because it received activation from the perceptual input. Thus, once a functional relation has been recognized and encoded, it is no longer needed to simulate the transformations which led to its discovery.

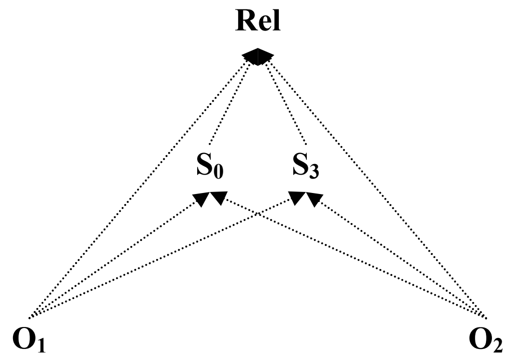

Figure 5.2.7. Representation of a relation. There are two bundles of structures (states) which define the initial description of the relation between the objects (S0) and the state after several transformations (S3). In order to encode S3 one has to foresee the transformations of the initial state by simulating the executions of actions.

The dynamics of relational categorization presented in Figure 5.2.8 highlights several predictions of the embodiment approach to relational representations. First, it makes clear that the recognition of certain relations (such as functional relations) is dependent on how well and how quickly one can simulate a series of actions and predict their consequences. Hence, the recognition of relations could be constrained by the ability of the human body to execute or simulate a series of actions in quick succession. This prediction is in concordance with the results of several studies (e.g. Bhalla & Proffit, 1999; Proffitt, Stefanucci, Banton & Epstein, 2003; Flusberg, Jenkins & Boroditsky, 2009) which have shown that the discovery and encoding of relations could be constrained by physical characteristics of the human body and its ability to execute certain actions. On the other hand, the experiments of Dixon & Dohn (2003), Trudeau & Dixon (2007) and Day & Goldstone (2009) provide indirect evidence for the hypothesis that relational categorization is dependent on the ability to predict the consequences of simulated actions

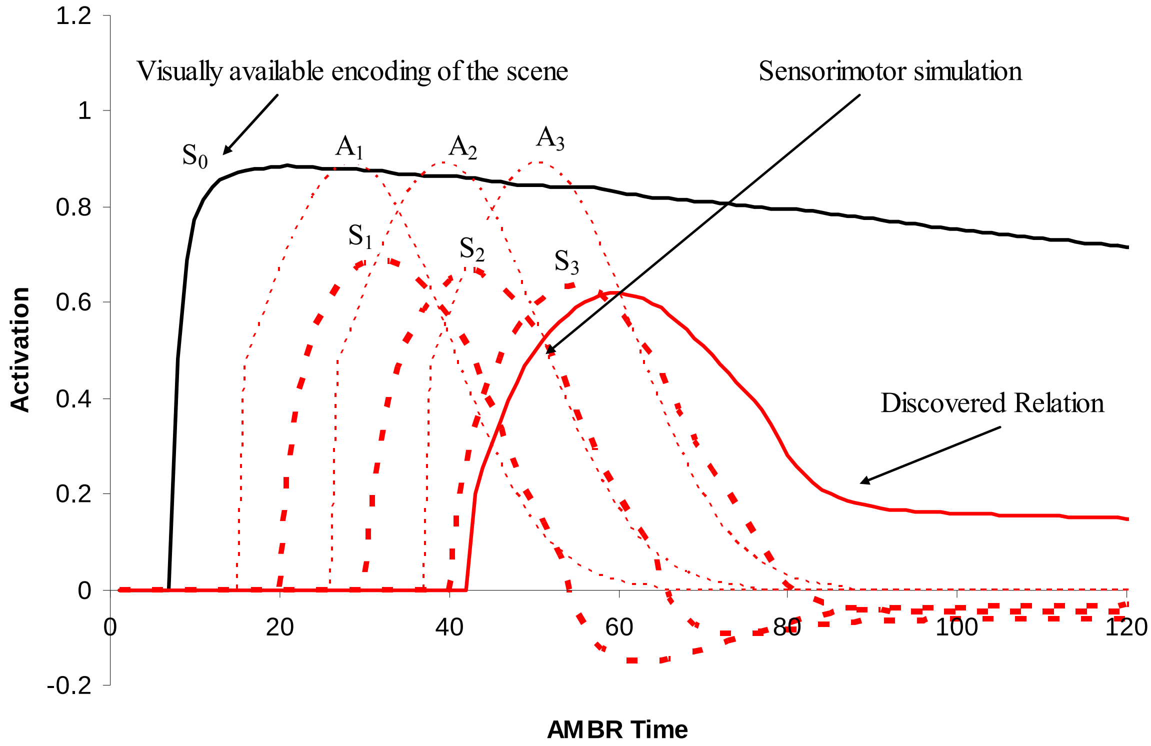

Figure 5.2.8. Dynamics of relational categorization. The states representing simulated transformations of the initial scene (S_1, S_2, S_3) appear after actions (A_1, A_2, A_3) are simulated and gradually fade away. The red line represents the activation of the relation which has been discovered.

Another prediction of the embodiment approach which is highlighted by Simulation 2 is concerned with relation comparison. Apparently, in order to compare two relations one must first discover them. Figure 5.2.9 depicts the dynamics of discovering two relations in succession and mapping their representations. Two bundles of simulated transformations were needed to recognize the relations. The mapping of the represenations of the discovered relations started only after both simulations were completed. Hence, the model predicts that relation comparison is dependent on the ability to run the two perceptual simulations in close temporal proximity. The relation comparison will be hindered if the actions underlying the two simulations are incompatible and it is hard to execute them in succession. For example, the model predicts that people will perform slower when comparing relations which require the same body effector to be used for simulating different actions.

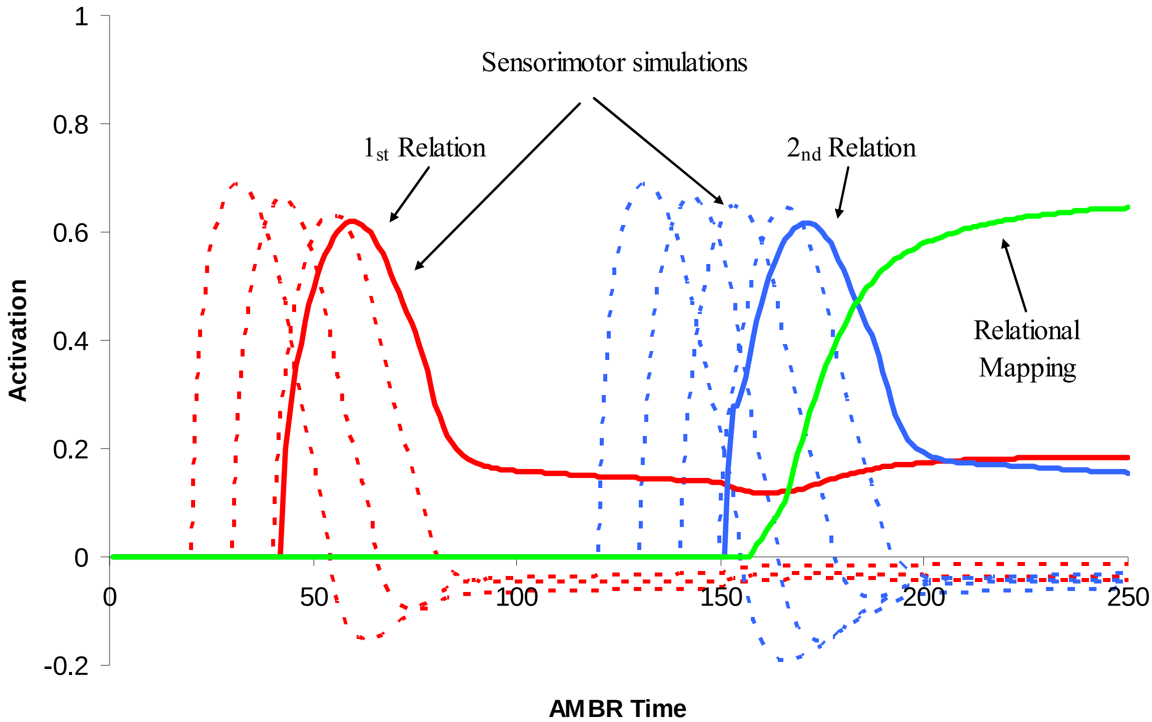

Figure 5.2.9. Dynamics of relation comparison. The green line represents the degree of mapping of the two relations (the activation of the hypothesis that maps them). The mapping of the relations starts as soon as both relation instances are available. The interval between thee two perceptual simulation reflects the constraints of the body to run two perceptual simulations in close temporal proximity.

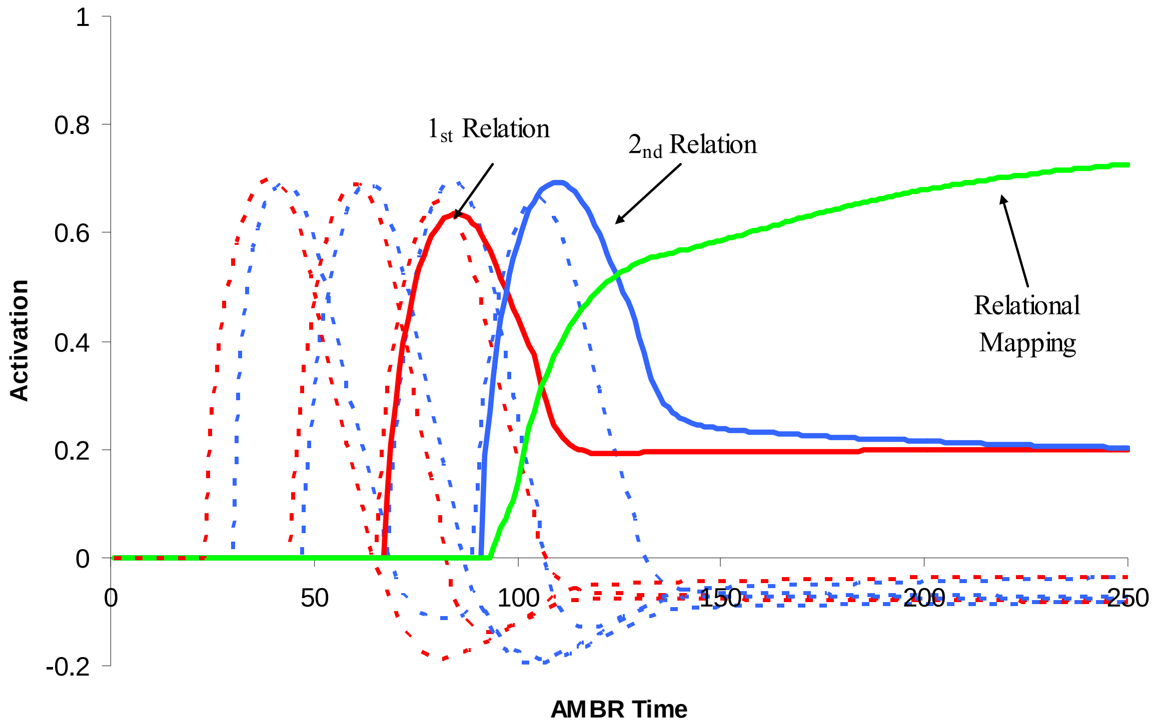

Figure 5.2.10. Parallel processing of relations. The two sensorimotor simulations are run simultaneously and the encoding of the relations proceeds at the same time. Thus the mapping of the relations starts much earlier than when relations are processed sequentially (Figure 5.2.9).

Recent studies have shown that priming procedural knowledge may facilitate responses in presumably purely conceptual tasks (Myung, Blumstein & Sedivy, 2006; Day & Goldstone, 2009). Our model of grounding relations in simulated interactions with the environment predicts similar effects. First, we predict that the processing of a relation will be facilitated if we prime the procedural knowledge which is involved in its simulation. Second, we predict that relations will be compared faster if they involve the simulation of similar actions. Such an effect will occur because the representation of an already simulated and encoded relation would be activated during the simulation of the other relation. Action similarities can also have a more direct effect if we adopt a parallel view of simulating and comparing relations.

We have assumed so far that the comparison of relations is a sequential process – the relations are encoded one after the other and then their representations are compared. However the proposed model does not preclude the possibility that the relations are encoded and compared *in parallel* if the nature of the simulated actions admits their simultaneous execution. Figure 5.2.10 demonstrates such a scenario. It is obvious that the comparison of relations proceeds much faster when the interactions underlying the relations are simulated in parallel.

The parallel comparison of relations has several benefits apart from maximally reducing the time interval separating subsequent sensorimotor simulations. If the transformations underlying relational meaning are simulated in parallel it is possible to compare not only the final states of the transformations, but also the internal states, as well as the motor programs which are involved in the simulations. Such comparisons would facilitate the mapping of relations given that the two simulations share common patterns of perceptual and motor information.

The last prediction of the model that we would like to highlight is that the parallel processing of relations would be most efficient when it is possible to *dynamically align* the perceptual simulations underlying relational meaning. Dynamical aligning means that corresponding states of two simulated transformations should occur in approximately the same moments. The effect of dynamical aligning would be particularly strong if the simulations have perceptual and motor commonalities, but it would facilitate relation comparison even if only the final states can be mapped.

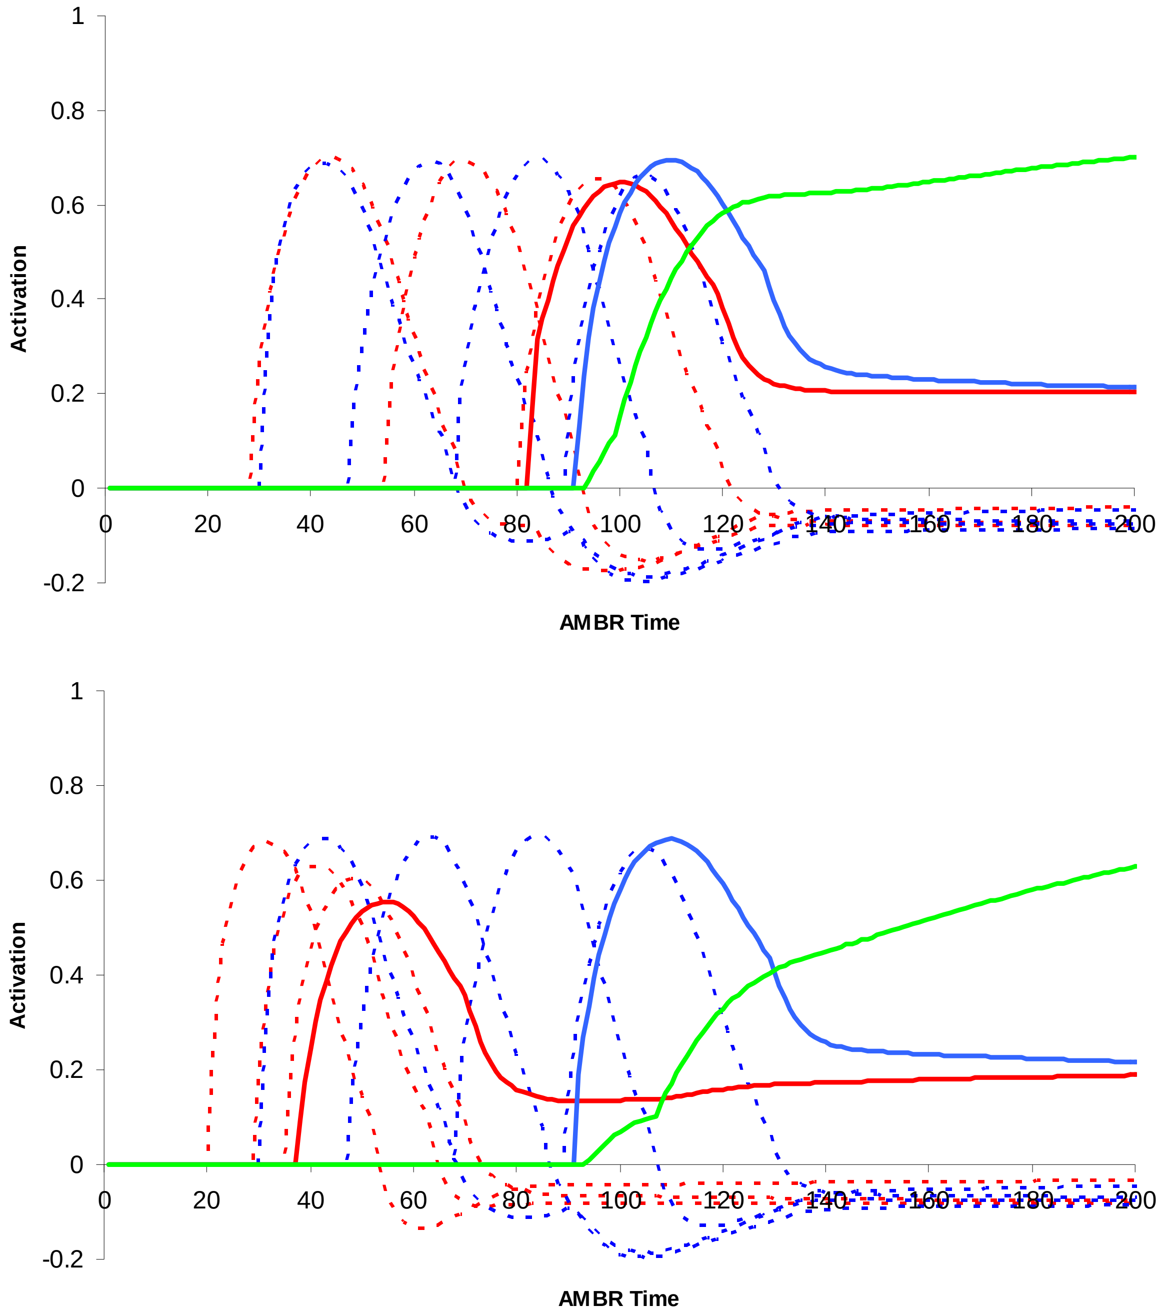

Figure 5.2.11. The advantage of dynamically aligning two sensorimotor simulations. The two simulations start at the same, but in one of the cases (top) the simulations are dynamically aligned. When the simulations are not dynamically aligned (bottom) the peaks of the recognition hypotheses do not coincide in time and the mapping of the relations is slower (the slope of the green line is flatter).

Figure 5.2.11 shows two cases in each which the simulations of relations are either dynamically aligned or not. We can see that relation comparison is facilitated when the two relations are recognized and encoded at approximately the same time. This happens because the effect of simulating the execution of actions gradually fades away and thus the representations of relations are most active immediately after the simulation (look at the curves of the thick red and blue lines). In order to efficiently compare such representations, one has to simulate them in such a way, that their peaks of activation coincide in time. The figure also demonstrates that in the case of a parallel processing of relations, the synchronization of simulated actions is more important than the total simulation time (i.e. the sum of the two simulations times).

The prediction that relation comparison is enhanced by dynamically aligning sensorimotor simulations could have a variety of manifestations. For example, it asserts that people would compare faster relations involving actions which can be coordinated. Also, we predict that an asymmetry of the human body, such as hand dominance, would have an effect on comparing relations to the extent that more pronounced hand dominance would impede the dynamic aligning of simulated actions which involve both hands.

Discussion

Simulation 2 demonstrated 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 sensorimotor 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 sensorimotor 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

This chapter 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.

Experiment 1a

The goal of the first experiment we conducted was to provide initial evidence that the representations of relations involve sensorimotor interactions with the environment. If indeed perceiving and reasoning with relations involve physical execution or simulations of actions, then these processes should be affected by the constraints of the body which executes or simulates these actions.

One way to investigate the role of the body in relational representations is to exploit the idea of affordances. Recent studies in the stimulus response compatibility paradigm have found evidence for the intrusion of task-irrelevant motor representations during a perceptual judgement (e.g. Tucker & Ellis, 1998, 2004; Richardson, Spivey, & Cheung, 2001; Symes, Ellis & Tucker, 2005, 2007; Borghi et al, 2007; Ellis, Tucker, Symes & Vainio, 2007; Vingerhoets, Vandamme & Vercammen, 2009; Bub & Masson, 2010; Jax & Buxbaum, 2010). The specifics of the interactions between perceptual and motor information 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.

Experiment 1 aimed to provide evidence that the affordances of the objects participating in a given relation influence the way the relation is processed. To this end, a task was devised which required subjects to compare the relations between two pairs of objects with varying affordances.

The proposed view of relations postulates that the functional relations between objects are grounded in physically executed or simulated interactions with the objects. Therefore, it is predicted that the affordances of the objects would constrain the process of comparing relations. Subject would compare relations faster when it is easier to them to manipulate the objects.

A second prediction was that participants would solve the task more efficiently if they were able to simulate the interactions underlying the two relations in close temporal proximity (see Simulation 2 for a justification of this prediction). Hence, we expected that there would be not only main effects 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 in close temporal proximity or at the same time.

Method

Paricipants 40 right-handed participants (24 females) took part in the experiment for course credit or as volunteers. Their average age was 23.15 years (age range from 17 to 34, *SD* = 3.19).

Stimuli The stimulus set was constructed out of 144 photos of various household objects. A stimulus consisted of two pairs of objects. One of the objects in each pair was tool-like and manipulatable (e.g., a fork, a hammer, a saw). The objects were paired in such a way that there would be a functional relation between them, for example 'hammer' – 'nail' (the relation is 'used-to-drive-in'), 'key' – 'lock' ('unlocks'), 'fork' – 'spaghetti' ('used-to-eat-with'), etc. We will refer to the relations between the objects displayed in the left and the right part of the screen as 'left relation' and 'right relation'. The left and the right relations were analogous in the target stimuli (Figure 6.1.1a) and non-analogous in the fillers (Figure 6.1.1b). A pre-test study was used to select the objects participating in the target stimuli in such a way that there was maximal agreement among people that the relations were analogous. The objects in each relation were oriented in such a way, that the intended interaction with them would be more convenient for the left or for the right hand of the participant. We will say that a relation had a left affordance when it was easier to grasp and manipulate the participating objects with the left hand and we will say it had a right affordance when it was easier to grasp and manipulate the objects with the right hand - Figure 6.1.2.

Figure 6.1.1a An example of a target stimulus. The relations in the target stimulus are analogous ('X-is-used-to-cut-Y'). The left relation has a left affordance - it easier to grasp the axe with the left hand and. The affordance of the right relation is right because it is easier to grasp the chopper with the right hand.

Figure 6.1.1b An example of a filler. The relations in the target stimulus are not analogous. Both affordances are left.

All images were resized to 400x400 pixels. The two relations were separated by a vertical line. The tool-like objects were always displayed at the bottom position.

Design The experiment had a 2x2 within subject design. Two independent variables were defined by the affordances of two relations: *affordance of the left relation* and *affordance of the right relation*. We will refer to each experimental condition by a pair of letters. The affordance of the left relation is indicated by the first letter (L – left affordance, R – right affordance) and the affordance of the right relation is indicated by the second letter (Figure 6.1.2):

The dependent variable was the response time of verbal responses (‘yes’/’no’).

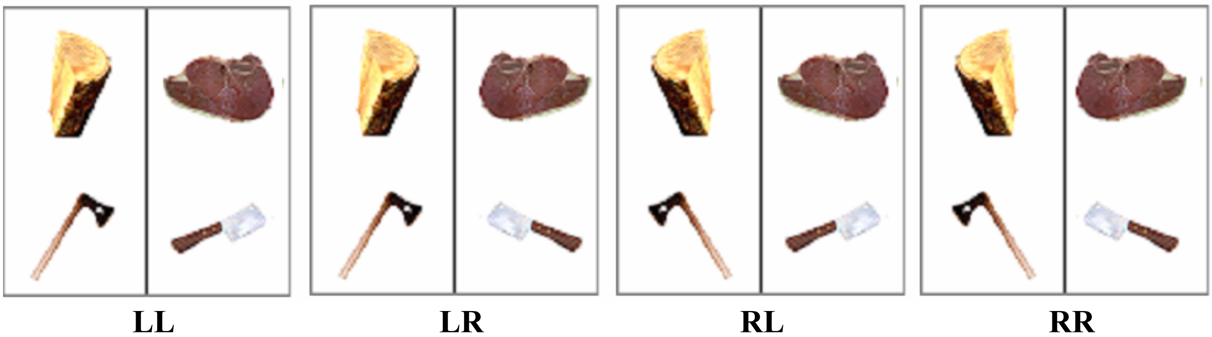

Figure 6.1.2 An example of the manipulation of the affordances of relations in Experiment 1. The first letter of the name of the experimental condition indicates the affordance of the first relation and the second letter indicates the affordance of the second relation. For example, in this example RL indicates that the axe is oriented in such way that it is easier grasp it with the right hand and the chopper is easier to be grasped with the left hand.

Procedure Each stimulus was presented once to each subject. The experimental conditions and the locations of presentation were counterbalanced across subjects. For example, half of the subjects saw the 'axe-log' relation presented in the left part of the screen and the other half saw it in the right part of the screen. The sequence of target and filler trials was pseudo-randomized. The stimuli were presented in the same order to all subjects.

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. During the training session, participants were presented with two simple geometric shapes (squares, triangles, stars, etc) and were asked to respond by saying 'yes' if the two shapes were identical and pronounce 'no' otherwise. If the experimental program was unable to detect the response, it warned the participant to respond louder and more clearly. All participants went through 20 such microphone training trials and were not admitted to do the rest of the experiment unless they had completed successfully at least 85% of them.

The real experimental session started with eight 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 was generated. Participants responded by saying 'yes' or 'no'. The inter-trial interval was 2500 ms. There was an additional task after one fourth of the trials (selected randomly). The additional task required the participants to respond whether a given word denoted any of the two relations just presented. The additional task was introduced in order to urge the subjects to comprehend the relations.

Stimulus presentation and response recordings were controlled by E-prime software (Schneider, Eschman, & Zuccolotto, 2002). The experimenter stayed with the subjects during the experiment and marked the response to each trial during the inter-trial interval. The experimenter also marked invalid trials in which the subject fail to articulate their response clearly. The experiment took about 15 minutes.

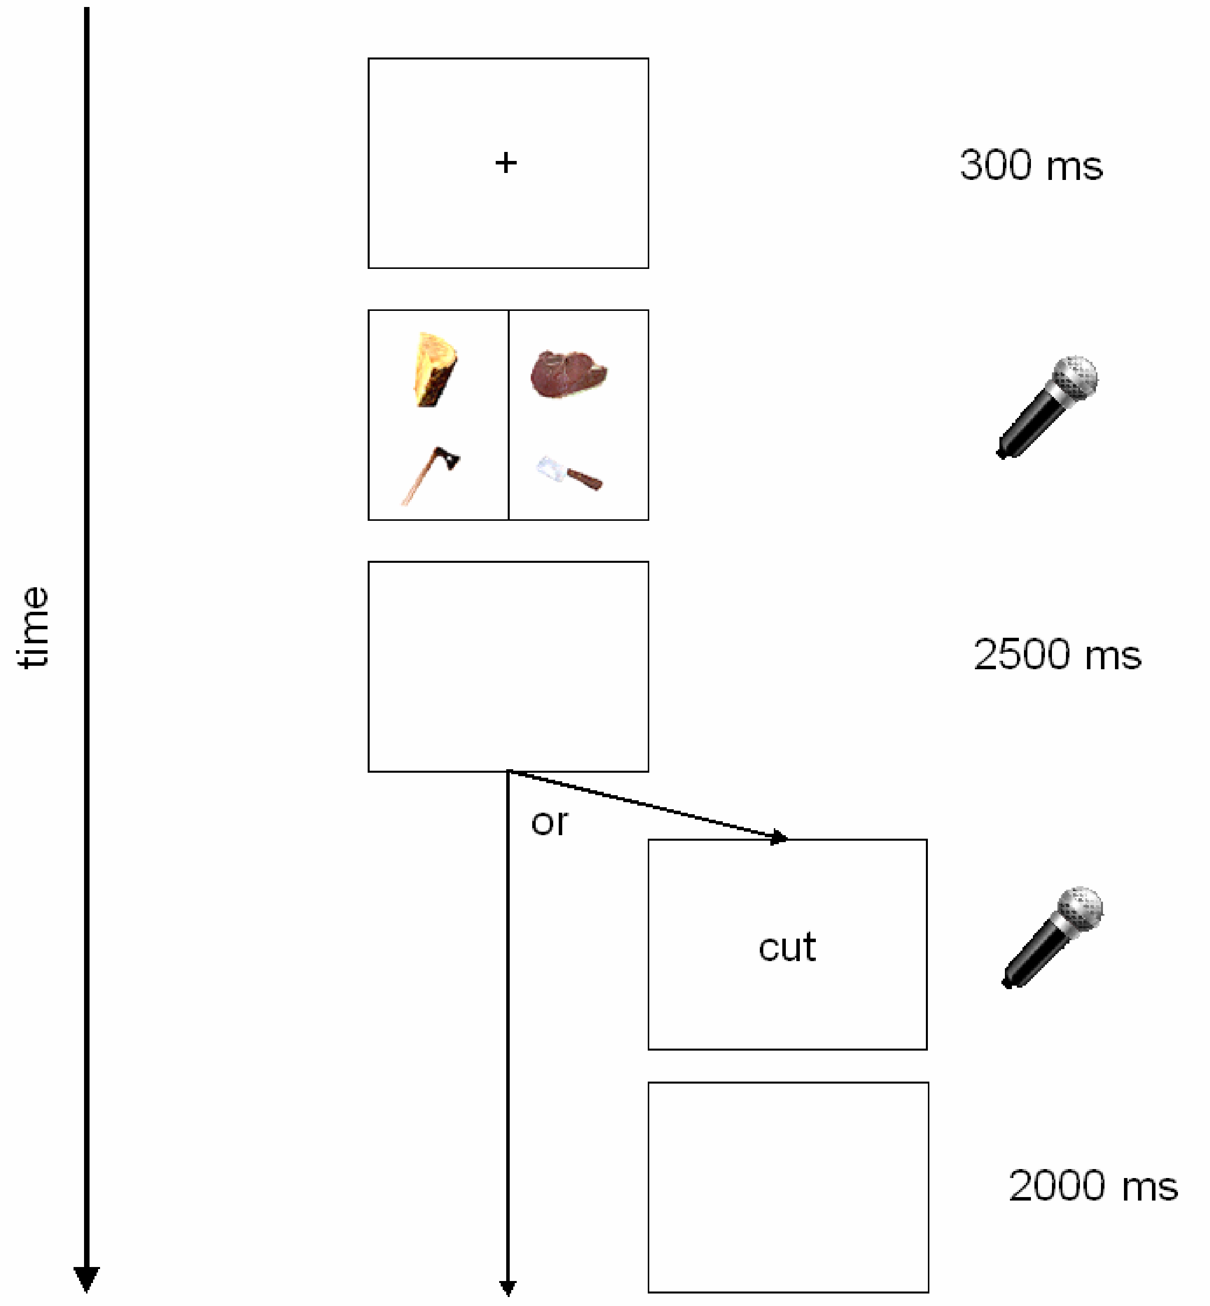

Figure 6.1.3. Experimental procedure of Experiment 1. After one fourth of the trials the subjects were asked to respond whether a given word denoted any of the relations.

Example	Experimental Condition	RT (ms)	SD (ms)
	**LL**	1833.78 ± 114.04	356.59
	**LR**	1786.30 ± 112.65	352.22
	**RL**	1908.28 ± 104.61	327.10
	**RR**	1876.97 ± 122.49	383.01

Table 6.1.1. Mean subject response times, confidence intervals and standard deviations obtained in the four experimental conditions of Experiment 1a.

Results

Only correct 'yes' responses were included in the analysis of response times. An incorrect response was counted if the subject responded by saying 'no' in a target trial (in which the relations were analogous). Incorrect responses constituted 6% of the non-filler data. Responses times lying more than ±2.5 standard deviations from the RT mean were excluded from the analysis. Thus, a total of 92.80% of the originally collected non-filler RT data were included in further analysis. There was no significant difference in the number of removed data points between the four experimental conditions (χ^2(3, N = 70) = 5.77, *p* = .12).

Table 6.1.1 presents the subject response times means data in each of the four experimental conditions. People responded fastest when the affordance of the left relation was left and the affordance of the right relation was right.

A 2x2 (affordance of left relation x affordance of right relation) repeated measures ANOVA was performed on subject RT means and revealed a statistically significant main effect of the affordance of the left relation ($F(1, 39) = 8.72$, $p < .01$, *ES* = .18) but failed to find an effect of the affordance of the right relation ($F(1, 39) = 1.98$, $p = .17$, *ES* = .05) – Figure 6.1.4. The interaction between the two factors was not significant ($F(1, 39) = .05$, $p = .82$, *ES* = .00).

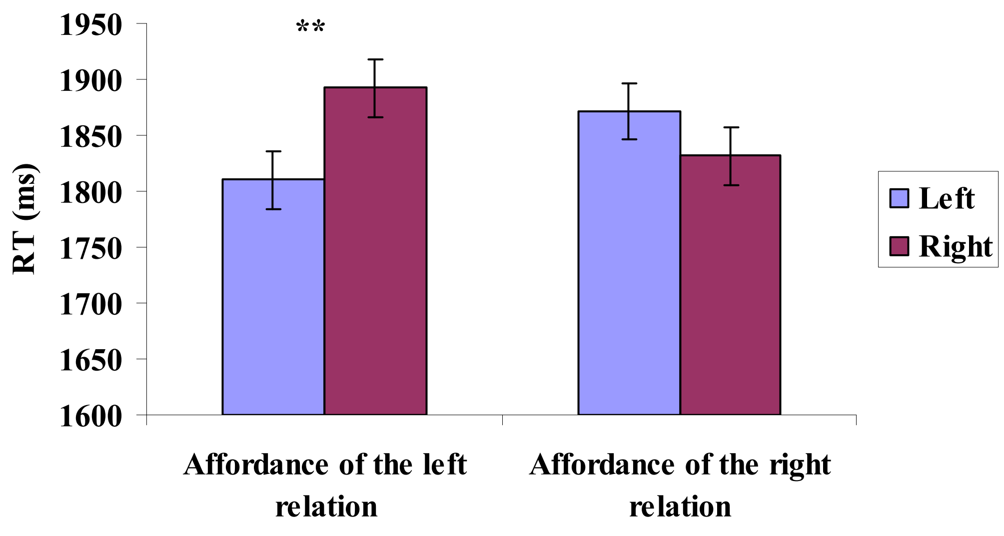

Figure 6.1.4. Mean subject response times with two affordance factors. The size of the effect of the affordance of the left relation was bigger. All participants were right-handed. The error bars indicate standard errors.

Discussion

The shortest response times in Experiment 1 were in the case when the affordances of the relations were congruent to the nearest hand. In other words, people responded faster when the objects displayed in the left part of the screen were easier to be manipulated with the left hand and the object displayed to the right were easier to be manipulated with the right hand. Interestingly, the size of the effect of the affordance of the relation displayed near the subject's non-dominant right hand was bigger.

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 affordance of the relation which was presented near the subject's dominant hand. Other studies have shown that the affordance effect is stronger for the dominant hand (e.g. Richardson, Spivey, & Cheung, 2001). However the result is easy to explain if we assume that people tried to use both their hands in simulating the actions involved in the two relations. If this is the case then the relation on the non-dominant side becomes the critical one, as in this case the subjects have to simulate an action in an unusual way, using their non-dominant hand, while their dominant hand is engaged in another action.

Several researchers have shown that mere looking at tool-like objects activates regions of the brain related to action (e.g. Beauchamp et al., 2002; Beauchamp & Martin, 2007; Buccino et al., 2009). One may argue that the effects of Experiment 1 could be due to such automatic activation of motor information and not to simulating interactions with the objects. This motor information could enable the recognition of tool-like objects. Hence people could just be faster to recognize objects which are easier to be grasped. Such an explanation does not render support for the hypothesis that the perception and comparison of relations involve simulating interactions between objects. Therefore we decided to conduct a complementary experiment.

Experiment 1b

The goal of this experiment was to make sure that the findings of Experiment 1 were specific to the task of relation comparison. We devised another task, in which people had to compare the same tool-like objects with varying affordances, but the task did not involve perceiving and comparing functional relations.

Method

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 tool-like objects with varting affordances which were used in the target stimmuli of the previous experiment. People saw two objects per screen. We refer to the object displayed in the left part of the screen as 'left object' and will use 'right object' to refer to the object displayed in the right part of the screen.

Design The design involved two independent variables: the affordance of left object and the affordance of the right object. The dependent variable was verbal response time. Only correct 'yes' responses were considered for analysis.

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 18 filler trials was compiled using 18 photos of man-made objects, none of which was used in the target trials, and other 18 photos of natural objects (fruits, plants, rocks, etc).

Figure 6.1.5. Example stimuli used in Experiment 1b. The correct response is 'yes' for the stimulus on the left (none of the objects is of natural origin) and 'no' for the stimulus on the right (one of the objects is of natural origin). Only the correct 'yes' responses were analyzed. The pictures of the objects were the same that were used in the target trials in Experiment 1a.

Results

Only correct 'yes' responses were included in the analysis of response times. Responses times lying more than ±2.5 standard deviations from the RT mean were removed. Table 6.1.2 presents the subject response time means data.

A 2x2 (left x right object affordance) repeated measures ANOVA was performed on subject means response times. It revealed a significant main effect of the affordance of the right objects (F(1, 15) = 4.80, p < .05, *ES* = .24) and failed to find an effect of the affordance of the left objects (F(1, 15) = .98, p = .34, *ES* = .06). The interaction was not significant (F(1, 15) = .20, p = .66, *ES* = .01). The results are presented in Figure 6.1.6.

Example	**Experimental Condition**	**RT (ms)**	**SD (ms)**
	**LL**	930.76 ± 83.79	157.24
	**LR**	903.99 ± 83.58	156.85
	**RL**	962.71 ± 86.67	162.64
	**RR**	911.09 ± 85.71	160.85

Table 6.1.2. Mean subject response times, confidence intervals and standard deviations obtained in the four experimental conditions of Experiment 1b.

An analysis of mean item RTs yielded the same pattern of results. The effect of the affordance of the left object was significant (F(1, 17) = .45, p = .51, *ES* = .03) and the effect of the affordance of the right object was significant (F(1, 17) = 5.49, p < .05, *ES* = .24). The interaction was not significant (F(1, 17) = .44, p = .52, *ES* = .03).

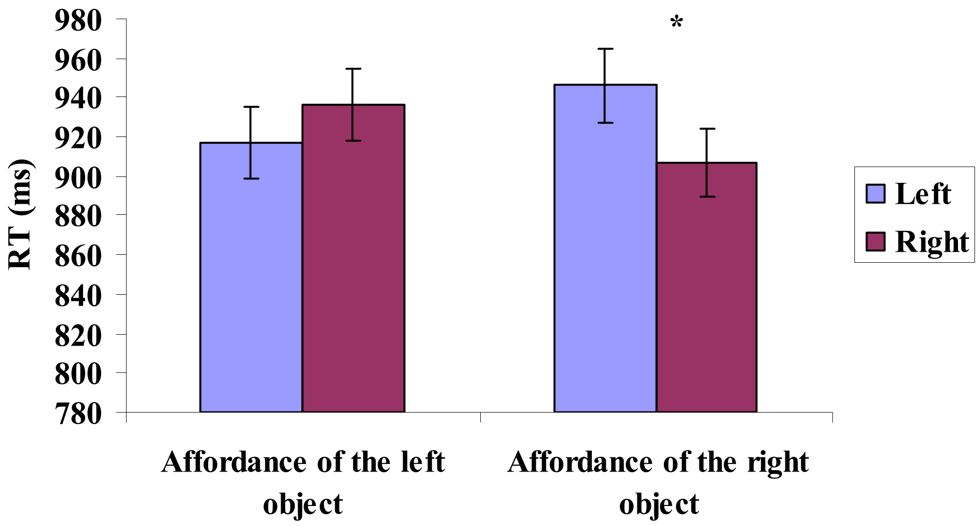

Figure 6.1.6. Mean subject response times in Experiment 1b. People responded faster when the affordances of the objects matched the positions of the objects – left affordance in left part of the screen and right affordance in right part of the screen. The size of the effect of the affordance of the right object was bigger than the size of the effect of the left object. The error bars indicate standard errors.

Discussion

Experiment 1b revealed that the effect of the affordances of objects displayed near the dominant hand is bigger. This result is different from what was found in Experiment 1a. We can therefore rule out the possibility that the main finding of the Experiment 1a – the bigger effect size of the affordance of the relation near the non-dominant hand - can be explained by automatic activation of motor representations which is not specific to task of relation comparison.

General Discussion of Experiment 1a and Experiment 1b

Experiment 1 provided initial evidence in support of the hypothesis that at least certain functional relations are embodied. First, it was shown that the affordances of objects constrain performance in a relation comparison task. Second, it was demonstrated that the effect of the object affordances can not 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, 2009). If the activation of motor areas was just a side effect it would not have had any effect on pronouncing the word ‘yes’ which was the way participants indicated their response in the target trials. Informal debriefing after the experiments showed that subjects were completely unaware that the task had anything to do with their hands and simulations of actions.

Experiment 2a

Experiment 1a managed to find evidence that object affordances constrain the process of comparing relations. The major finding was that the overall response time was affected mostly by the affordance of the relation displayed near the subjects' non-dominant hand. We reasoned that this effect occurred because subjects attempted to use both their hands in order to simulate the two relations in close temporal proximity.

However the design of Experiment 1 did not allow to control the sequence in which the subjects look at the two relations. Therefore it is possible that the effect of the relation which had been attended last was different (bigger or lesser) from the other one. A new experiment was designed and conducted in order to overcome this problem.

The experiment used the same stimuli as Experiment 1a, but the relations 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 elimination of the factor of presentation location served to set apart the effect of the affordances of the stimuli from any spatial compatibility effects. It is well known that people respond faster to stimuli which location is compatible to the response action (Simon & Rudell, 1967). Although the response action in Experiment 1a did not include any spatial codes, it is possible that subjects' response times had been affected by the congruence of the presentation location and the affordances of the stimuli. For example, an interaction between objects with a left affordance could be easier to be simulated with the right hand if they are displayed in the right part of the screen.

The new design also allowed testing the effect of a stimulus – the relation which was presented first – which had to be retrieved from memory at the time of the comparison. If any affordance effect was found for the first relation, it would seriously question any disembodied view of relational comparison which assumes that relations are first encoded as symbols and then compared.

Method

Participants 36 right-handed participants (20 females) took part in the experiment for course credit or as volunteers. Their average age was 24.06 years (age range from 18 to 53, SD = 5.91).

Stimuli The stimulus set was the same one that was used in Experiment 1a. We will refer to the relaiton which appears first as 'first relation' and to the relation which is presented second as 'second relation'.

Design The experiment had a 2x2 within subject design. The two independent variables were: *affordance of the first relation* and *affordance of the second relation.* We will refer to the experimental conditions by a pair of letters: the first letter indicates the affordance of the first relation and the second letter indicates the affordance of the second relation. For example, 'LR' refers to the experimental condition in which the affordance of the first relation is left and the affordance of the second relation is right.

The dependent variable was the response time of participants' verbal 'yes'/'no' responses.

Only correct 'yes' responses were considered for analysis.

Procedure Each stimulus was presented once to each subject. The experimental conditions and the order of presentation of the relations (first or second) were counterbalanced across subjects. It was made sure that the same experimental condition would not repeat more than 3 times in a row. Trials with target stimuli and fillers were pseudo-randomized, so that a given correct response would not repeat more than 3 times. The trial sequence was fixed for all subjects, i.e. they saw the stimuli in the same order.

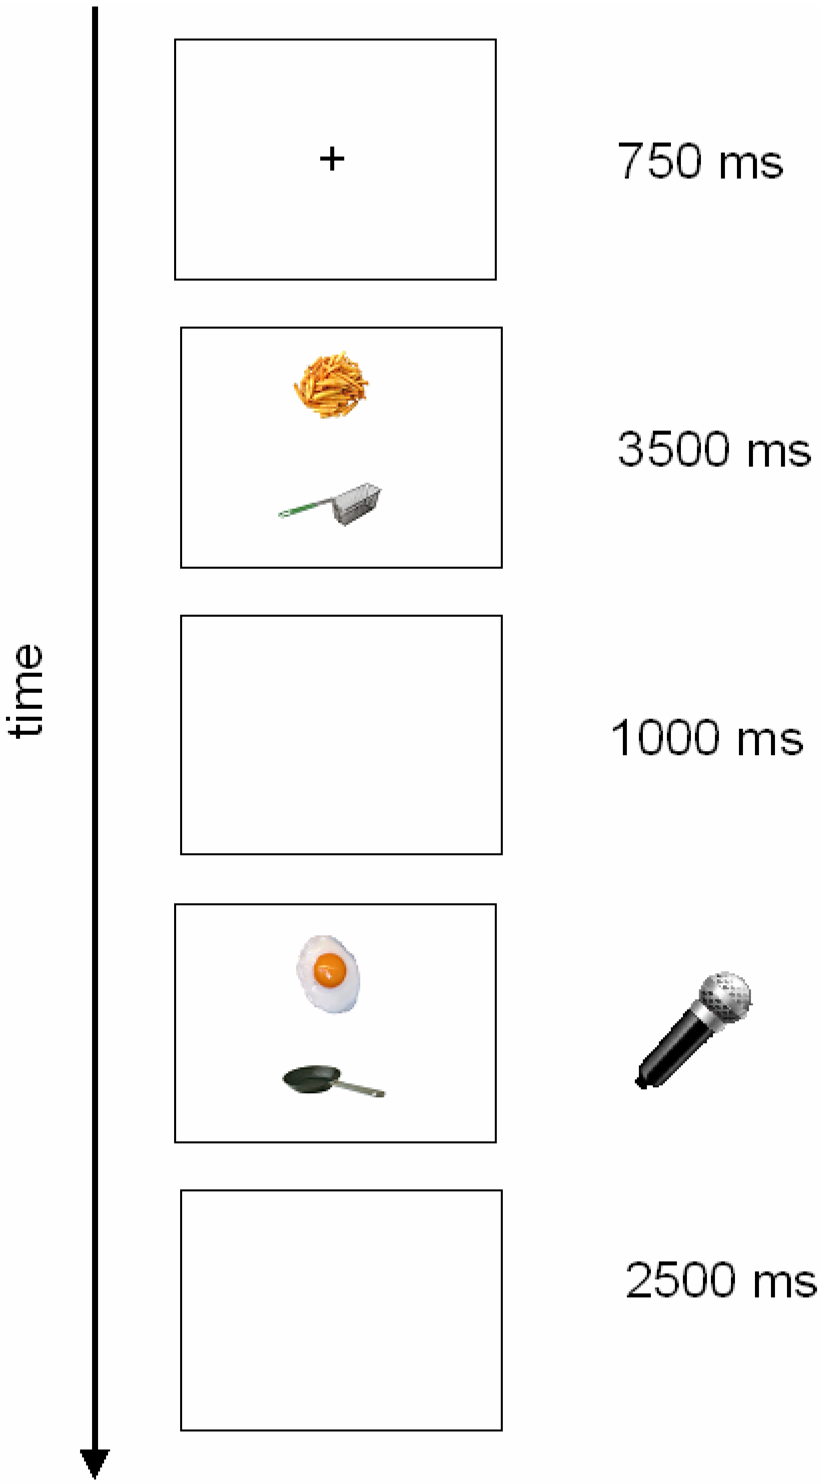

Figure 6.2.1. Experimental procedure of Experiment 2a. Response time was measured from the onset of the second relation until a verbal response was given.

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 would 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 (750ms), followed by the onset of the first pair of objects. The objects were displayed one below the other in the centre of the screen. Subjects were instructed to perceive the relation between the objects without making any response. The first pair of objects was presented for 3500 ms and when it disappeared the screen stayed blank for 1000 ms. After that a second pair of objects was presented at the same position as the first one. The stimuli stayed on the screen for 5000ms or until a response was generated. Participants were instructed to respond by saying 'yes' if the second relation was analogous to the first one and say 'no' otherwise. The subject's response time (RT) was measured since the onset of the second relation till the moment a verbal response was detected. The inter-trial interval was 2500 ms. The total number of test trials for each subject was 36, including 18 target trials and 18 filler trials. The experiment took about 15 minutes.

Results

Only correct 'yes' responses were included in the analysis of response times. The number of incorrect responses constituted 9.6% of the non-filler data. Trials in which subjects failed to respond were excluded from the analysis. RTs lying more than ±2.0 standard deviations from the RT mean time were also removed. Thus a total of 86.30% of the originally collected non-filler RT data were included in the analysis. There was no significant difference in the number of removed data points between the four experimental conditions ($\chi^2(3, N = 82) = 1.61, p = .66$).

The subject mean response time data are presented in Table 6.2.1. The shortest reaction times were in condition LR in which the affordance of the first relation was left and the affordance of the second relation was right.

Example	**Experimental Condition**	**RT (ms)**	**SD (ms)**
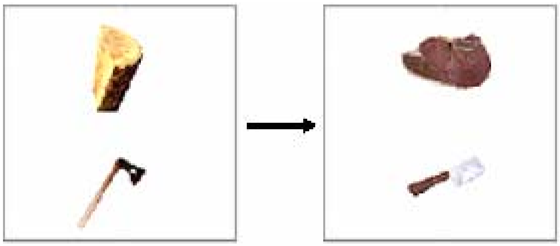	**LL**	1205.09 ± 87.95	259.93
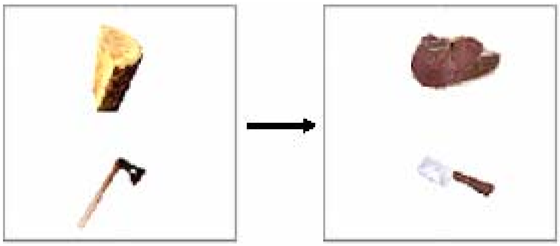	**LR**	1182.72 ± 87.12	257.48
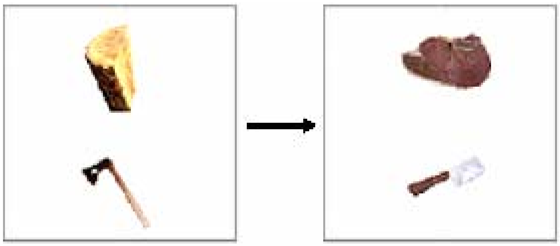	**RL**	1294.00 ± 93.13	275.24
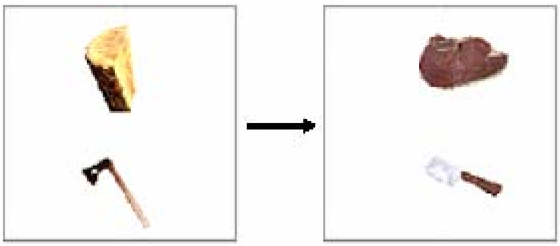	**RR**	1243.61 ± 102.28	302.29

Table 6.2.1. Mean subject response times, confidence intervals and standard deviations obtained in the four experimental conditions of Experiment 2a.

A 2x2 repeated measures ANOVA was performed on subject RT means and revealed a significant main effect of the affordance of the first relation ($F(1, 35) = 7.12$, $p < .05$, $ES = .17$). There was no significant effect of the affordance of the second relation ($F(1, 35) = 2.02$,

p = .16, ES = .06). The interaction between the two affordance factors was not significant ($F(1, 35) = .20$, $p = .66$, $ES = .01$). The results are presented in Figure 6.2.2.

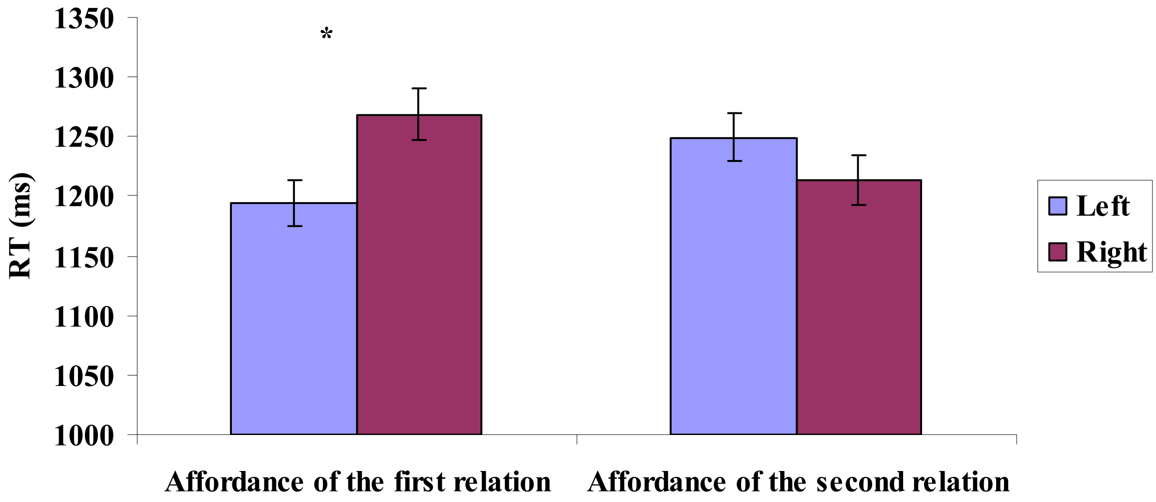

Figure 6.2.2. Experiment 2a results. Subjects' responses were significantly faster when the affordance of the first relation was left. The tendency for the affordance of the second relation was reversed. The error bars indicate standard errors.

An analysis of mean item response times also found a main effect of the first affordance ($F(1, 17) = 9.05$, $p < .01$, $ES = .35$). The effect of the second affordance ($F(1, 17) = 2.69$, $p = .12$, $ES = .14$) and the interaction ($F(1, 17) = 2.15$, $p = .16$, $ES = .11$) were not significant .

Discussion

The results replicated the findings of Experiment 1a as long as we found an effect of the object affordances of the relations. Also, the shortest response times were in the condition when one of the affordances was left and the other one was right. Another similarity is that the effect size of the affordance of the first relation was bigger than the effect size of the affordance of the second relation.

The novel finding of Experiment 2a is that right-handed subjects were faster to respond when the affordance of the first relation was left. Unlike in Experiment 1a, this result can not be explained by presentation location as all stimuli were presented in the centre of the screen. At first glance, there is no reasonable explanation why participants would be faster when the affordance of one of the relations is congruent to their non-dominant hand.

However the results start to seem logical if we make an additional assumption. Let us assume that subjects always simulated the interactions between the objects in the second relations with their dominant hand (i.e. the right one – all participants were right-handed). Such an assumption is not unreasonable – people usually use their dominant hand to grasp objects, especially tools, unless they are doing anything else with it or unless the affordance of the object is such that it is much easier to grasp it with the non-dominant hand (which is not the case when the object is presented in the centre of the screen). If people simulated the interactions between the objects of the second relations with their dominant hand then they had no chance but to use their non-dominant hand for simulating the first relation when having to compare the two relations. Such a scenario also explains why we found any effect of the affordance of the first relation – the first relation was retrieved from memory and re-simulated when subjects saw the second relation and tried to compare the two relations. Richardson, Spivey, & Cheung (2001) have demonstrated that the affordances of objects retrieved from memory might have an effect on the current task.

The pattern of results of Experiment 2a is inconsistent with any classical symbolic account of relational reasoning. If relations are first encoded as symbolic structures and then compared, then there would not be any effect of the first affordance. The first relation would have already been encoded by the time the second relation is presented and the response is given.

The different directions of the affordance effects of the first and second relations indicate that people attempted to use both their hands in order to simulate the two relations in close temporal proximity. It the affordance effects were due to automatic activation of motor information which facilitated object recognition than the directions of the two affordance factors should have been the same. Nevertheless, we conducted a complementary experiment to make sure the results of Experiment 2s are specific to the relation comparison task.

Experiment 2b

This experiment aimed to show that the results of Experiment 2a were specific to the task of comparing relations. The stimulus set and design and procedure were almost identical to Experiment 1b, except that the objects were displayed one after the other.

Method

Participants 24 right-handed participants (17 females) took part in the experiment for course credit or as volunteers. Their average age was 22.79 years (age range from 17 to 32, SD = 3.13).

Stimuli The stimuli were the same as in Experiment 1b, but the objects were presented sequentially. We will refer to the object which was presented first as ‘first object’ and use the term ‘second object’ for the object which appeared second.

Design The design was the same as in Experiment 1b. The affordances of the objects were described by two independent variables – *affordance of the first object* and *affordance of the second object*. The dependent variable was verbal response time. Only correct ‘yes’ responses were considered for analysis.

Procedure The setting of the experiment was similar to Experiment 2a. Each trial began by a fixation cross (750 ms), followed by the presentation of the first object (2000ms). After that, the screen stayed blank for 1000 ms and the second object was presented. Subjects were instructed to say ‘yes’ if none of the objects was of natural origin and say ‘no’ otherwise. Response time was recorded since the onset of the second object. All objects were displayed in the centre of the screen. The order of presentation of the objects and the experimental conditions were counter-balanced across subjects.

Example	**Experimental Condition**	**RT (ms)**	**SD (ms)**
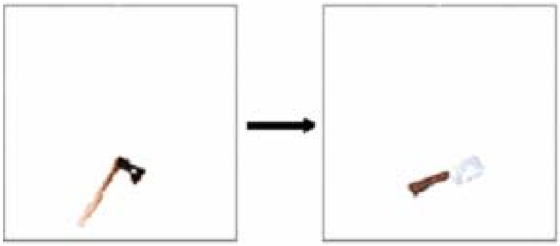	**LL**	785.50 ± 64.90	153.69
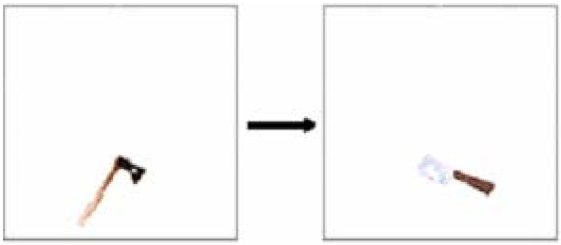	**LR**	751.02 ± 55.17	130.65
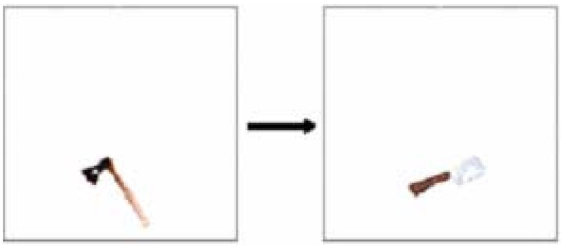	**RL**	750.00 ± 57.82	136.93
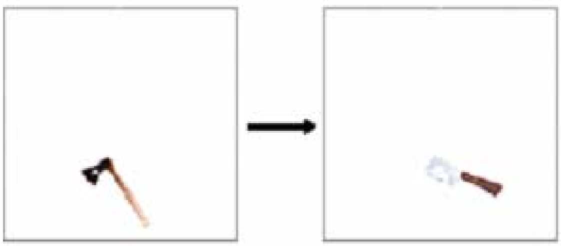	**RR**	732.14 ± 57.07	135.16

Table 6.2.2. Mean subject response times, confidence intervals and standard deviations obtained in the four experimental conditions of Experiment 2b.

Results

Only correct 'yes' responses were analyzed. Invalid or incorrect responses were excluded from the analysis. Response times lying more than ±2.0 standard deviations from the mean RT time were removed. Thus a total of 92.40% of the originally collected non-filler RT data were included in further analysis. Table 6.2.2 presents the subject response time means data.

A 2x2 repeated measures ANOVA was performed on subject mean RTs. It revealed main effects of the affordances of the first ($F(1, 23) = 5.18$, $p < .05$, *ES* = .18) and the second object ($F(1, 23) = 5.36$, $p < .05$, *ES* = .19). The interaction was not significant ($F(1, 23) = 0.22$, $p = .64$, *ES* = .01). Response times were faster when the affordances of both objects were right (Figure 6.2.3).

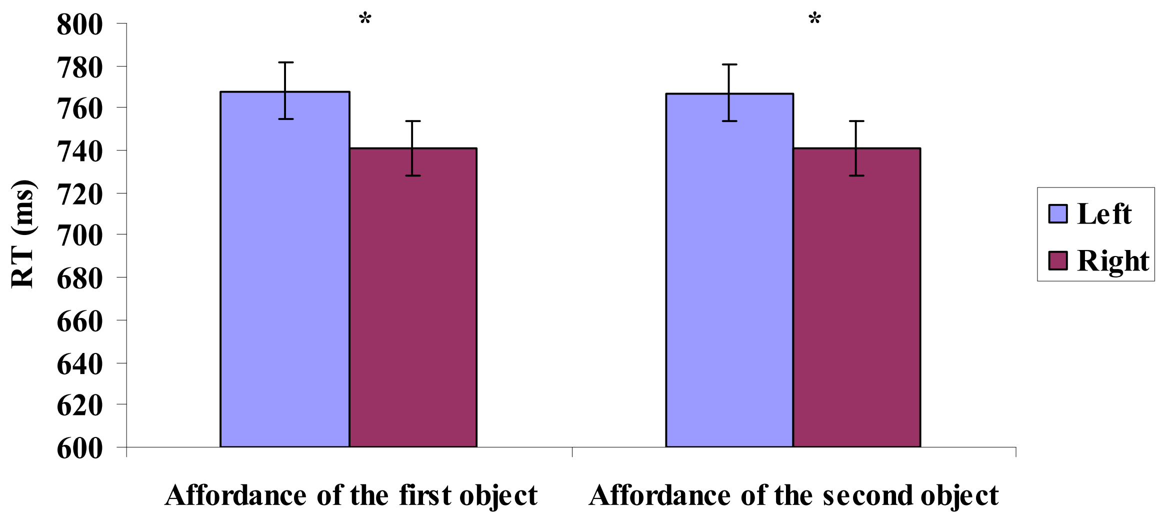

Figure 6.2.3: Results in Experiment 2b. Right-handed subjects were significantly faster when the affordances of both objects were right. Such a result is consistent with the hypothesis that the recognition of tool-like objects involves simulating interactions with them, but cannot explain the results of Experiment 2a. The error bars indicate standard errors.

Discussion

Experiment 2b showed that the findings of Experiment 2a can not be explained only by processes related to object recognition. Response times were shorter when the affordances of both objects were congruent to the subject's dominant hand. Also, there was no difference between the sizes of the effects of the affordances of the objects. These results are different from what was found in the previous experiment and they show that the results of Experiment 2a are specific to the process of comparing relations.

General Discussion of Experiment 2a and Experiment 2b

Experiment 2a provided further evidence in support of the hypothesis that the meaning of functional relations is grounded in the sensory-motor dynamics resulting from simulated interactions with the environment. The pattern of results is also consistent with the idea that comparing functional relations involves running two or more such simulations simultaneously or in close temporal proximity. The outcome of Experiment 2b rules out the possibility that the results were due to the object affordances per se.

The outcomes of the experiments are clearly in support of an embodied view of relations. However one may attempt to interpret the results of Experiment 2a without adopting the specific idea of embodying relational representations and relational reasoning by referring to the theory of event coding (Hommel, Müsseler, Aschersleben & Prinz, 2001). According to this theory, elements of perception and action are encoded in a common medium. When the stimulus features related to perception and action are active for a long time period they become bound into an event file. Once bound, these features are less available for planning of other actions. Hence it is possible that a right affordance of the first pair relation would bind the features representing the right hand of the subject to the stimulus features of the first relation. When the second relation is presented, the right hand of the subject would be less available for simulating the use of the presented objects and the response would be delayed. A result of this kind has been reported by Richardson, Spivey, & Cheung (2001). Such an explanation reduces the role of simulated action to the process of object recognition.

However, the theory of event coding contradicts the results of the Experiment 2b, unless it is assumed that the presentations times were too short for the event filing to happen. Such an assumption is highly unlikely to be true, as in the complementary experiment the first object was presented for a fixed period of 2000ms, followed by a 1000 ms inter-stimulus interval before the second object was displayed. This period is much longer than the time which was required for suppression of future actions in the studies of Spivey, Richardson & Cheing (2001) and Stoet & Hommel (1999). Also, there is no evidence so far that such a phenomenon could occur outside the stimulus-response compatibility paradigm and have an effect on the timing of the response action which was used in our experiments. Hence, the event filing explanation can not adequately account for the results obtained by Experiments 2a and 2b.

Experiment 3

Experiments 1a and 2a managed to provide evidence that people simulate interactions between objects when comparing functional relations. Moreover, the results of both experiments imply that the process of functional relation comparison is constrained by characteristics of the human body such as hand-dominance. Experiment 1a showed that the effect size of the affordance of the relation near the non-dominant hand is bigger. It was speculated that this asymmetry occurred because people attempted to use both their hands in order to simulate the relations in close temporal proximity. Experiment 2a found another asymmetry – right handed subjects were faster to compare sequentially presented relations when the affordance of the relation was left and the affordance of the second relation was right. Experiments 1b and 2b ruled out the possibility that these effects were due to processes related to object recognition but they are specific to the process of relation comparison.

Simulation 2, described in Chapter 5, demonstrated that relation comparison would be facilitated if it is possible to run the two underlying perceptual-motor simulations in parallel and dynamically align them. In order to this one has to be able to synchronize the start and ends points of the simulations. Such synchronization would be difficult if asymmetric limbs are involved in the two simulations.

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. To this end, we decided to directly induce such asymmetries by attaching weights to the arms of the subjects. Two predictions were formulated. First, it was predicted that attaching weights to the arms of the subjects would be reflected in the response times in a relation comparison task. Any such effect would provide additional evidence that people do engage in simulated sensorimotor interaction when comparing relations. The second prediction, which was specific to this experiment, was that subject would respond faster when the manipulations of their body was symmetric. The rationale was that an asymmetric manipulation would lead to differentiation of the dynamics of the simulations of the relations which are compared and thus the comparison process would be impeded.

Method

Paricipants 64 right-handed participants (32 females) took part in the experiment for course credit or as volunteers. Their average age was 23.83 years (age range from 18 to 55, *SD* = 5.17).

Stimuli The stimulus set was the same as in Experiment 1a.

Figure 6.3.1. The affordances of the relations were not manipulated in this experiment. The affordances were always congruent to the nearest hand: the affordance of the relation was always left and the affordance of the right relation was right.

Design The affordances of the relations were not manipulated – the stimuli were always presented as in the LR condition from Experiment 1 (Figure 6.3.1). The only independent variable was defined by the weights attached to the wrists of the participants – the *weights* factor. Thus the following experimental conditions were defined:

L – The weight was attached to the left wrist of the subject.

R – The weight was attached to the right wrist of the subject.

LR – There were weights attached to both wrists of the subject.

No – There were no weights on the subject's wrists.

The design of the experiment was between-subject - each subject participated in one experimental condition only. The dependent variable was verbal response time. Only correct 'yes' responses were considered for analysis.

Procedure At the beginning of the experiment each subject was assigned to one of the four experimental conditions. Subjects were then instructed how to attach the weight(s) to their wrists(s) (except for the 'No' condition in which no weight was attached). The weights were textile strips full of heavy material (Figure 6.3.2). Each of them weighted 500g. The weights were firmly attached to subjects' wrists so that the subjects were able to move freely their arms.

Figure 6.3.2. The sports weights[1] which were used for manipulating the subjects' bodies.

[1] http://www.alibaba.com/product-gs/286216239/for_Nintendo_Wii_Fit_Ankle_Wrist_Weights.html

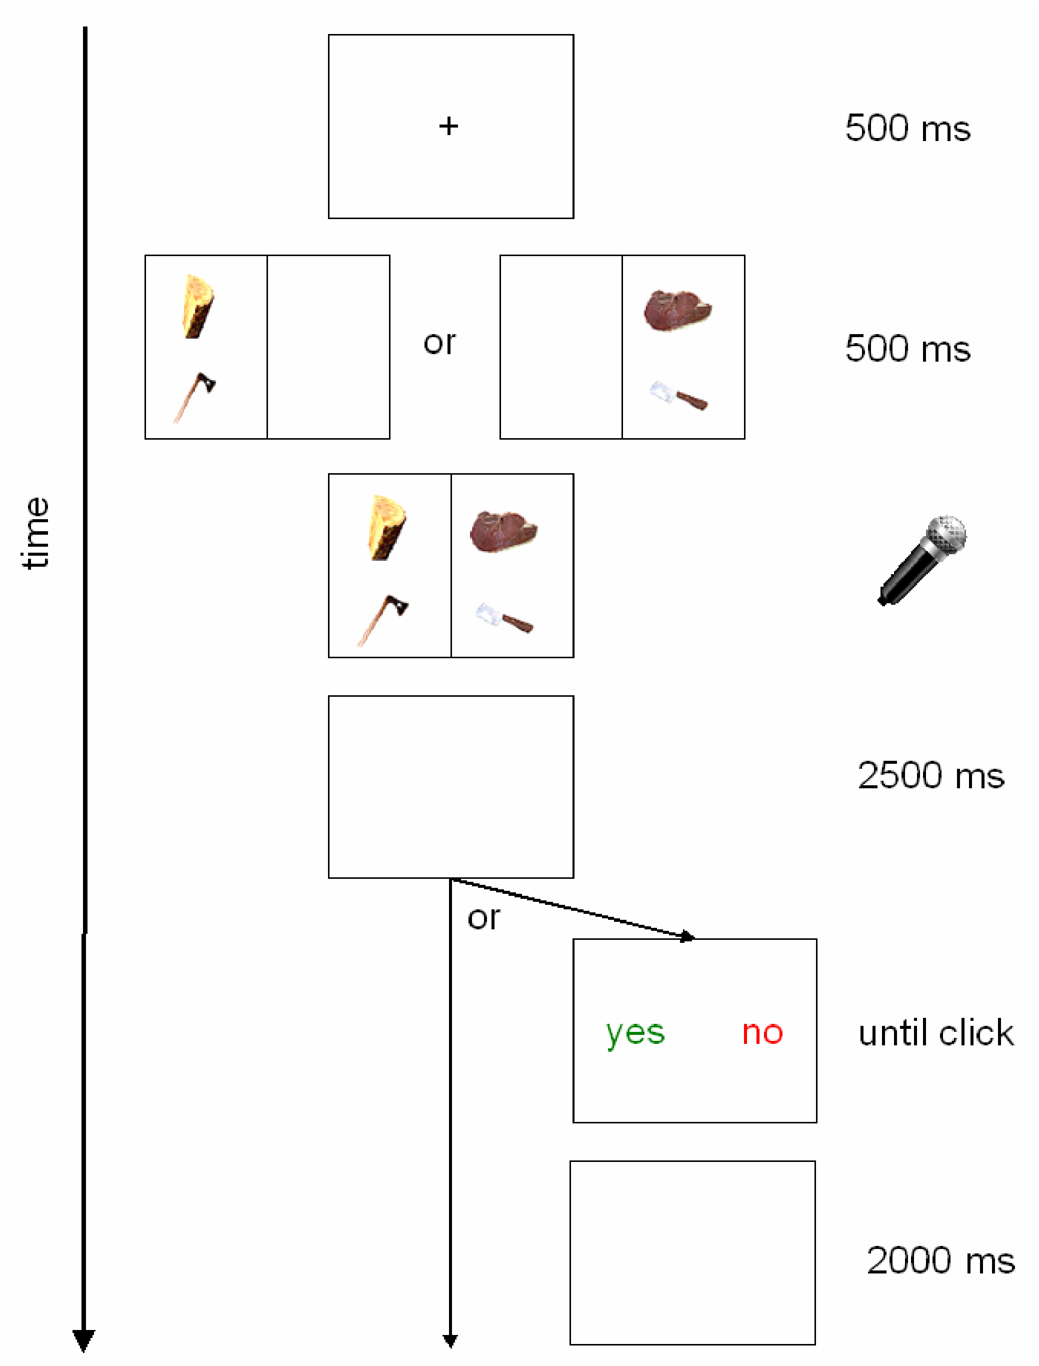

Figure 6.3.3. Experimental procedure of Experiment 3. One of the relations was displayed 500 ms before the other. After some of the trials the subjects were asked to reconfirm their response by clicking on the corresponding word. They were told to use their left hand when the word they wanted to click was in the left part of the screen and use their right hand if they wanted to select the word to the right.

The subjects were seated in front of computer with a 17” touch screen display and asked to play with it for a while in order to see how it works. The experimenter asked them to open a calculator application by using the touch screen display and perform a series of ten predefined arithmetic operations such as addition and multiplication of numbers. All subjects performed half of the operations with their left hand and the rest with their right hand.

When the familiarization with the touch screen display was completed the experiment continued in almost the same way as in Experiment 1a. There were just two slight changes to the procedure of Experiment 1a. First, in order to control for the order in which the subject look at the relations, the onset of one of the relations was delayed for 500ms. The order of presentation of relations was counter-balanced across subjects (for a given subject either the left or the right pair objects always appeared first). Response time was measured from the onset of the delayed relation. Second, in approximately one fourth of the trials the subjects were asked to reconfirm their response by pointing to the touch screen display and clicking on the ‘yes’ or ‘no’ words (Figure 6.3.3). This change to the procedure was made in order to keep the subjects aware of the weights attached to their wrists.

The experiment lasted about 15 minutes.

Results

Six subjects had more than 40% incorrect responses and their data were not analysed. Trials in which the remaining 58 subjects failed to articulate their response clearly were excluded from the analysis which resulted in removing 11.30% of the non-filler data.. Incorrect responses constituted 5.8% of the non-filler data. Only the remaining correct ‘yes’ responses were included in the analysis of response times.

The design of this experiment was between-subject. This was a problem as response times are subject to individual differences. In order in attempt to rule out the possibiility that any effects found are due to invididual differences, we searched for a way to normalize subjects' response times.

Example	**Experimental Condition**	**RT_{norm}**	**SD**
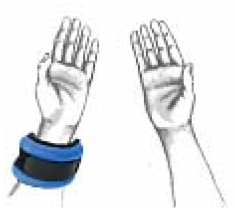	**L**	1.95 ± 0.17	0.33
	**R**	1.88 ± 0.25	0.51
	**LR**	1.70 ± 0.14	0.29
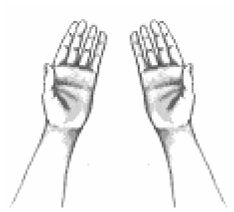	**No**	1.72 ± 0.16	0.32

Table 6.3.1. Mean subject normalized response times, confidence intervals and standard deviations obtained in the four experimental conditions of Experiment 3.

The procedure of the experiment required all participants to go through a microphone training section in order to make sure that they articulate their responses clearly enough (the description of the microphone training section is described in more detail in Experiment 1a). During the microphone training session the subjects had to respond whether two shapes are the same by saying 'yes' or 'no'. These responses were recorded and we were able to calculate the mean response time in the microphone training section for each subject. This value indicated how quickly a subject responds in a task which was not supposed to involve any perceptual-motor simulations. Hence the mean subjects' response time in the microphone session could serve as a baseline.

In order to normilize the response times of a subject in the actual experiment we divided them by the mean subject's RT from the microphone training session. In this way, we created a new depedent variable – normalized response time (RT_{norm}) –which allowed us to set apart the effects of the experimental conditions from the invidual differences between participants.

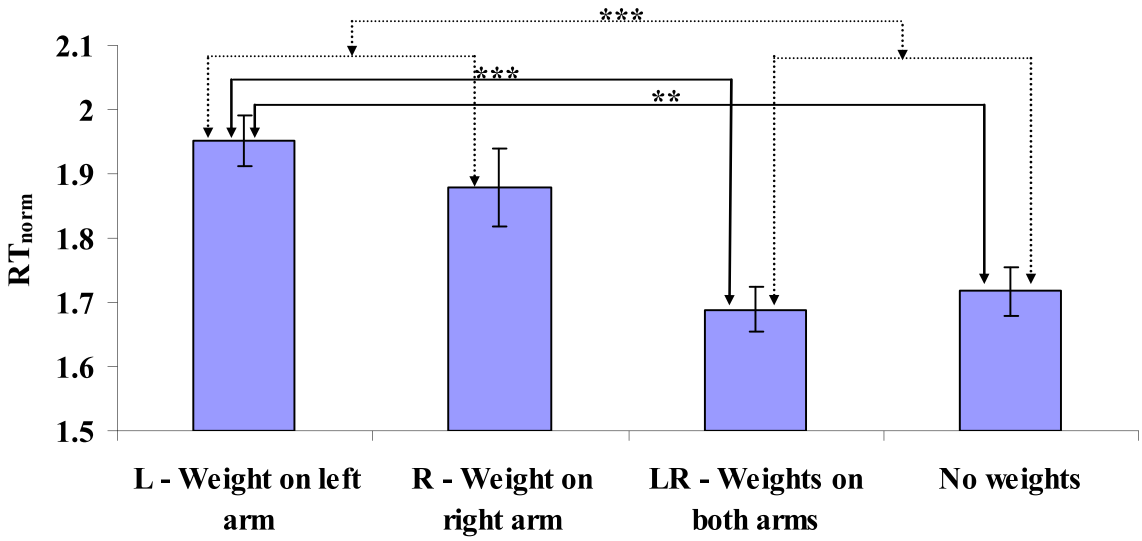

Figure 6.3.5 Mean item normalized response times. The longest response times occurred when a single weight was attached to the left wrist of the subjects (all subjects were right-handed). Error bars represent standard errors.

Normalized responses times lying more than ±2.5 standard deviations from the RT_{norm}

mean were excluded from the analysis which resulted in removing 2.30% of the correct responses.. A total of 80.40% of the originally collected non-filler response times data were included in further analysis. There was no significant difference in the number of removed data points between the four experimental conditions ($\chi^2(3, N = 205) = 1.03, p = .79$).

Table 6.3.1 presents the subject normalized response time means data in each of the four experimental conditions. People responded fastest in the LR and No conditions.

An one-way univariate ANOVA analysis of *subject* mean RT_{norm}s did not manage to find a significant effect of the weights factor ($F(3, 54) = 1.63$, $p = .19$, $ES = .08$). However, the analysis of *item* mean RT_{norm}s revealed a significant main effect ($F(3, 51) = 6.02$, $p < .001$, $ES = .26$). The results are presented in Figure 6.3.4. A set of Tukey HSD post-hoc tests revealed significant differences between conditions L and LR ($p < .001$) and between L and No ($p < .01$).

The major prediction that this experiment had to test was that people would respond faster when they their wrists were manipulated symmetrically. So we ran a contrast test which compared conditions the asymmetric conditions 'L' and 'R' against the symmetric conditions 'LR' and 'No'. A significant difference was found between symmetric and asymmetric conditions ($F(1, 17) = 14.91$, $p < .001$, $ES = .47$).

Discussion

The results of the experiment confirmed the two predictions. First, it was demonstrated that manipulations of subjects' bodies can have an effect in a task, which seemingly involves no limb-based motor activity. This result is consistent with the studies of Bhalla & Proffitt(1999) and Proffitt, Stefanucci, Banton & Epstein (2003) which also showed that manipulations of the physiological potential of the participants may affect their performance in a purely perceptual task. 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.

The fact that people responded faster when the manipulations of their arms were symmetric indicates that people compared not only the final results of the simulated interactions but also considered the dynamics of these interactions. Attaching a weight to one's wrist significantly changes the dynamics of the sensorimotor simulations which involve this limb. Therefore, when only one of the wrists is manipulated, it becomes harder to dynamically align the two simulations. Simulation 2 explicated the mechanisms that might underlie this comparison.

The analyses did not find any significant difference between the two symmetric conditions. People responded equally fast either when both their wrists were loaded with weights or none of them. This lack of difference implies that the effect of the weights factor was not related to the perception of the relations, but to the process of their comparison.

The analysis of the normalized response times revealed that attaching a weight to the left wrist of the subjects had a bigger effect than when the weight was on the right wrist. Such a result is reasonable considering the fact that all subjects were right-handed. When a weight is attached to the dominant hand of a subject, its movements are impeded and the dynamics of the simulated sensorimotor simulations related to this limb is made closer to the dynamics of the non-dominant one. On the other hand, when the movements of the non-dominant hand are impeded, the asymmetry becomes even bigger and the comparison of the relations becomes harder. This result is consistent with the results of Experiment 1 and 2, which revealed bigger effect sizes of the affordances of relations which were supposedly simulated with the non-dominant hand.

Experiment 4

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 a new experiment, 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 classical mental rotation task requires determining whether two images depict the same shape from different perspectives or that one of the shapes is a reflected version of the other. The perspectives of a shape are manipulated by rotating it in three (Shepard & Metzler, 1971) or two dimensional space (Cooper, 1975). It has been found the response time for 'yes' responses (when the shapes are the same) increases linearly as a function of the angle of rotation. The dominating explanation of this effect is that people mentally rotate one of the shapes until its visual image becomes identical to the other one. We have already discussed that there is evidence that transformations of mental images are at least in part guided by motor processes (Wexler, 1998; Wraga, 2003).

The mental rotation task provides an efficient way to manipulate the dynamics of the encoding and comparing relations. We devised an extended version of the mental rotation task – the dual mental rotation task - which required comparing the outcomes of two instances of the classical mental rotation task. The task involved comparison of relations in as long as determining whether two shapes are the same is a discovery of a relation between them ('sameness'). We predicted that people will respond faster when they are able to dynamically align the start and end points of the simulations involved in the two instances of the mental rotation task. In other words, performance in this task would depend on the absolute difference between the two rotation angles which define the durations of the mental simulations. People would respond fastest when the rotation angles were equal – in this case the two simulations would last the same time and it would be possible to dynamically align their start and end points.

Method

Paricipants 21 participants (12 females) took part in the experiment for course credit or as volunteers. Their average age was 23.95 years (age range from 19 to 36, *SD* = 4.36).

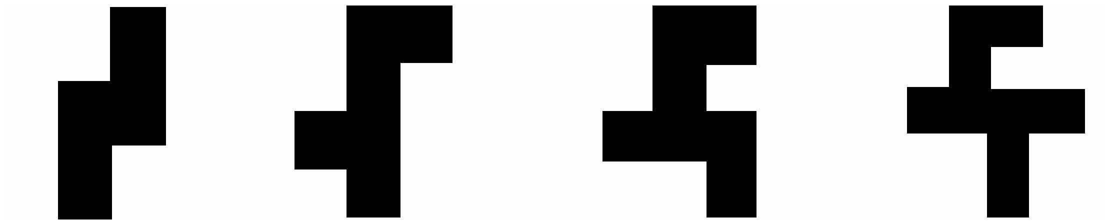

Figure 6.4.1a. The basic geometric shapes which used to construct the stimuli of Experiment 4.

Stimuli The materials were constructed out of 4 kinds of basic shapes (Figure 6.4.1). Each shape had 3 rotated variants (30°, 75°, 120°). Each stimulus consisted of two pairs of shapes. The first shape in each pair was one of the basic shapes. The second shape was a rotated version of the same basic shape or of its mirror copy. Thus, four kinds of stimuli were defined:

Same-Same: the shapes in the first pair were the same, as well as the shapes in the second pair.

Same-Different: the shapes in the first pair were the same and the shapes in the second pair were different.

Different-Same: the shapes in the first pair were different and the shapes in the second pair were the same.

Different-Different: the shapes in the first pair were different, as well as the shapes in the second pair.

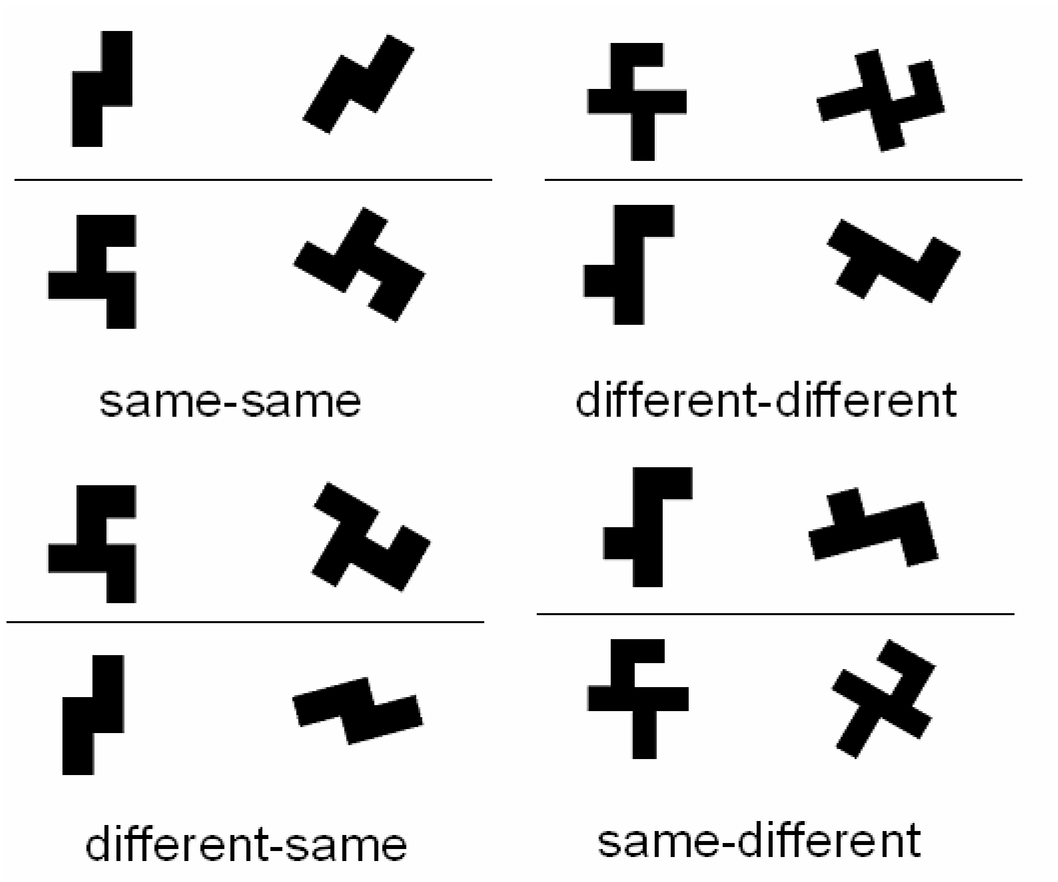

Figure 6.4.1b. Examples of the four kinds of stimuli used in Experiment 4. The top row contains examples of 'yes' trials because the correct response for these stimuli is 'yes' (the relations are the same). The bottom row demonstrates the two types of 'no' trials (the relations are different).

Design The independent variable was the absolute difference between the two rotations angles (A_{diff}). Only correct responses in '*same-same*' trials were considered for analysis.

The dependent variable was response time.

Procedure The stimuli set consisted of counterbalanced combinations of angles, combinations of basic shapes (Figure 6.4.1a) and stimuli kinds (Figure 6.4.1b). In order not to the make the experiment last too long the stimuli set was divided in 3 counterbalanced lists, each containing 144 stimuli. Participants were assigned to one of the lists in the beginning of the experiment.

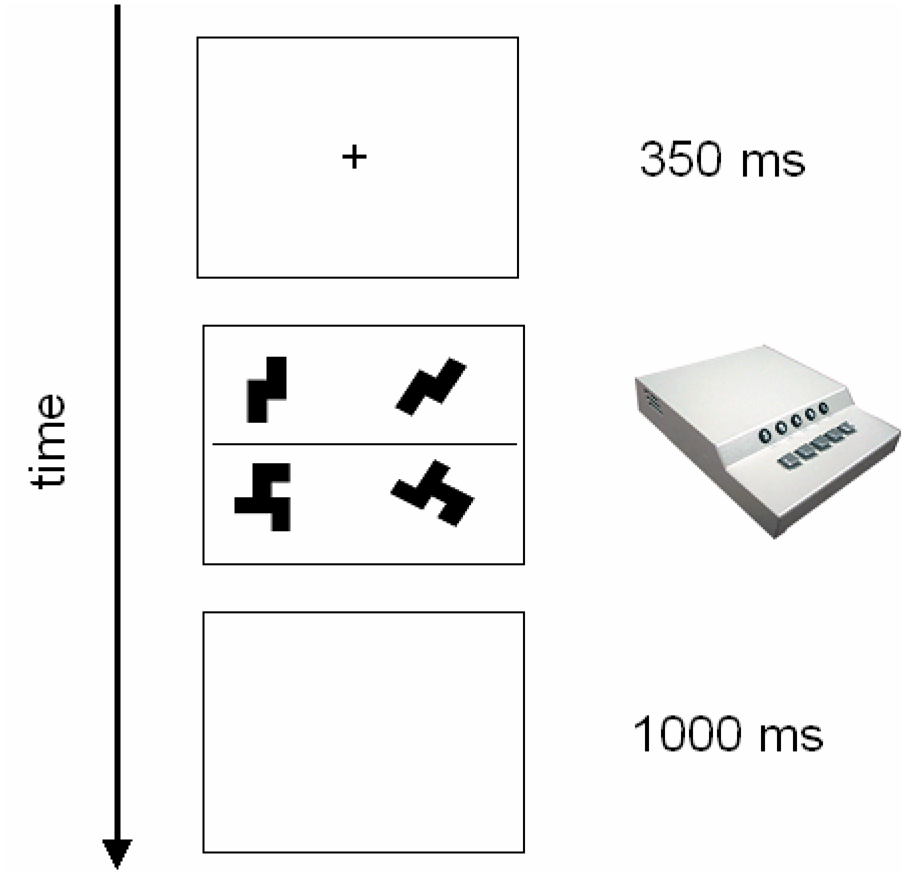

Figure 6.4.3. The experimental procedure of Experiment 4. Subject responded by pressing the first button of the serial button box if the relations in the two pairs of shapes were the same ('same-same' and 'different-different' trials) and by pressing the third button otherwise.

The experiment started with two practice sessions. First the participants did 10 practice trials of the classical mental rotation task. A pair of shapes appeared on the screen and the participants had to determine whether the two shapes were the same or not. After that the participants were introduced to the dual mental rotation task. It was explained that the task was to determine whether the relation in a pair of shapes is the same as the relation in another pair of shapes. The experimenter gave several examples of what 'same' and 'different' relations mean. Afterwards the participants did 16 practice trials of the dual mental rotation task. The angles of rotations in the practice trials were different from those used in the experimental trials (15°, 60°, 105°). There was feedback after each practice trial – the participants were informed whether their response was correct or not.

Each trial of the dual mental rotation task began with a fixation cross, which stayed on the screen for 350ms. After that two pairs of shapes were presented on the screen. The pairs were separated by a horizontal line (Figure 6.4.1b). The participants were instructed to respond as soon as possible by pressing predefined buttons of a button box. The inter-trial interval was 1000 ms. The procedure is shown in Figure 6.4.3.

Participants were tested in a sound-proof booth. The stimuli were being presented on a 19" computer monitor with a resolution of 640x480 pixels. Stimulus presentation and response recordings were controlled by E-prime software (Schneider, Eschman, & Zuccolotto, 2002). The experiment lasted about 20 minutes.

Results

Incorrect responses were excluded from the analysis fo response times. Responses times lying more than ±2.5 standard deviations from the mean response time were also removed. A total of 88.90% of the originally collected data from 'same-same' trails were included in further analysis. There was no significant difference in the number of removed data points between the three experimental conditions ($\chi^2(2, N = 84) = 1.14$, $p = .57$).

Example	**A_{diff} (degrees)**	**RT (ms)**	**SD (ms)**
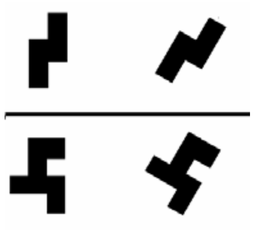	**0°**	2728.11 ± 323.46	710.61
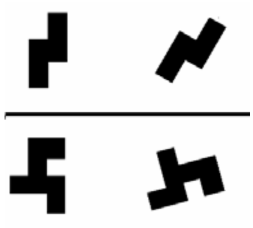	**45°**	2950.40 ± 333.35	732.32
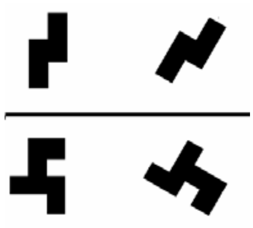	**90°**	3034.85 ± 295.50	649.18

Table 6.4.1. Mean subject response times, confidence intervals and standard deviations obtained in Experiment 4. The values of A_{diff} represent the absolute difference between the two angles of rotation.

The subject response time means data are presented in Table 6.4.1. A repeated measures ANOVA analysis of subjects' mean response time revealed a significant main effect of the A_{diff} factor ($F(2, 40) = 7.32$, $p < .01$, *ES* = .27). A contrast test revealed that the effect of A_{diff} is linear ($F(1, 20) = 12.86$, $p < .01$, *ES* = .39). The results are depicted in Figure 6.4.4.

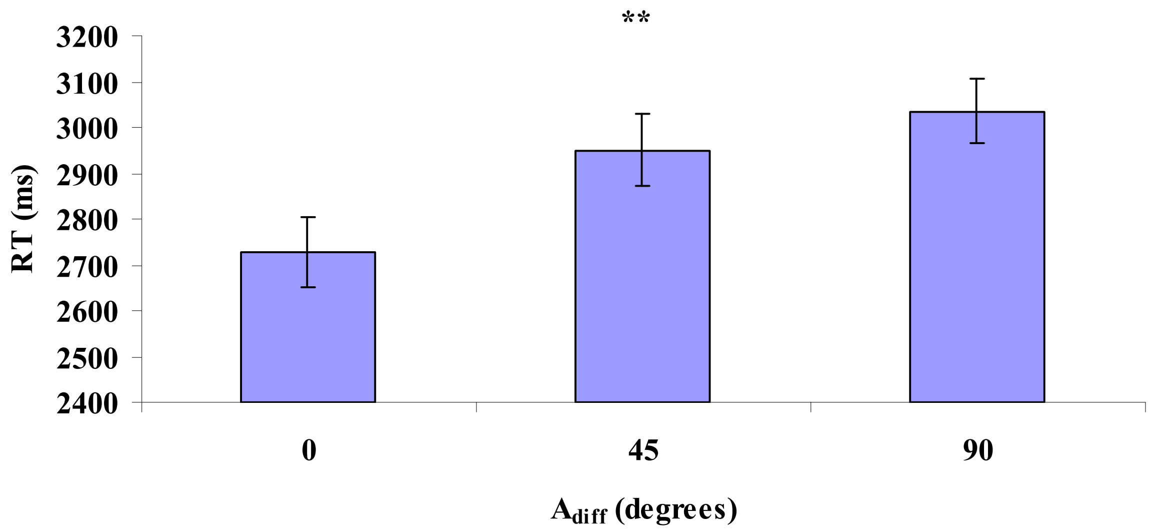

Figure 6.4.4. Mean subject response times for the different levels of the A_{diff} factor. The stimulus set was designed in such a way, that the mean total amount of rotation (A_1 + A_2) was the same in each of the conditions. The effect of A_{diif} on response times is linear. Error bars represent standard errors.

Discussion

The results of the experiment confirmed the prediction. People responded faster when the difference between the rotation angles was smaller.

The outcome of Experiment 4 renders support for the hypothesis that relation comparison can not be construed as a purely symbolic process. Kroger, Holyoak and Hummel (2004) found that manipulations of attributes and lower-level relations affect performance in comparing higher level relations. They concluded that people compared deep structures and not the just the relational symbols denoting the higher level relations. Experiment 4 extends this finding by showing that a quantitative manipulation of relational categorization can also influence the process of relation comparison.

The findings of Experiment 3 and Experiment 4 are similar. Experiment 3 showed that asymmetric manipulations of the human body lead to slowing down performance in comparing functional relations which supposedly involve simulated interactions with the environment. Experiment 4 demonstrated that people compare slower instances of the ‘sameness’ relation if they involve simulating actions with different dynamics. In both experiments responses were slowed down by inducing asymmetric changes to the dynamics of the simulated interactions underlying the relations. Both experiments also made certain that the effects were specific to the process of relation comparison. Thus we found converging evidence that temporal and motor dynamics plays important role in comparing relations.

Chapter 7. General discussion and conclusions

This chapter 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.

The benefits of grounding relations in action

A solution to the symbol grounding problem. 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 relational representations in action entails that they are dynamic as the execution or simulation of actions unfolds in time. Thus the embodiment account of relations is in tact with cotemporary theories and models of relational reasoning which converge to the understanding that there are no static features which define relational meaning.

A solution to the problem of context-sensitivity of relations. 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 (Glenberg, 1997; Barsalou, 1999).

A solution to the role-filler binding problem. The computational model of grounding spatial relations in action demonstrated by Simulation 1 outlines an elegant and efficient solution to the role-filler binding problem. It was shown that binding the arguments of relations to corresponding relational roles does not need introducing specialized mechanisms but could emerge from the dynamics of action execution. Although we did not thoroughly address the role-filler binding problem in this thesis, we believe that actions provide a natural and universal way of time locking the activation of relational arguments and relational roles. The action driven implementation of temporal role-filler binding is an alternative to the specialized mechanism suggested by other accounts of relational reasoning (e.g., Hummel & Holyoak, 1997), for which little empiric support has been found.

Falsifiability. 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.

Evidence in support of the predictions of the embodiment account of relations

Understanding relations involve perceptual-motor simulations. The results of Experiment 1, 2 and 3 supported the prediction that the recognition of relations involves running perceptual-motor simulations. Experiment 1 and 2 did this by showing that the affordances of objects affect verbal reaction times in a task requiring comparison of functional relations. Experiment 3 extended the findings of Experiment 1 and Experiment 2 by demonstrating by that a manipulation of the participants' bodies also affects performance in comparing function relations. The effects of object affordances and body manipulations can not be explained by traditional, disembodied, accounts of relational processing. The only feasible explanation is that subjects simulated grasping the objects and performing actions with them, as predicted by the embodiment account of relations.

Relation comparison is constrained by the ability of the body to execute two series of actions in close temporal proximity. Simulation 2 predicted that relation comparison would be faster if it is possible to run the two simulations in close temporal proximity, for example by employing different body effectors. Specific patterns of results found in Experiment 1, 2 and 3 support this prediction. In Experiment 1a we found a bigger effect size of the affordance of the objects near the non-dominant hand of the subjects. Experiment 1b showed that the result was specific to the task of relation comparison. Another interesting result emerged in Experiment 2a – right-handed subjects performed faster when the objects in the first relation were easier to be grasped with the left hand. Again Experiment 2b showed that this effect is specific to the relation comparison task. The results of Experiments 1 and 2 imply that participants attempted to use both their hands for simulating interactions between the objects. The results of Experiment 3 are consistent with this hypothesis. It was found that right-handed subjects responded slower when a weight was attached to their left arm, which indicates that they used both their hands for simulation. Generally, the results of Experiments 1, 2 and 3 are evidence that the process of relation comparison is constrained by the properties of the human body.

Relation comparison is enhanced by dynamically aligning perceptual motor simulations. Another prediction outlined by Simulation 2 was that the parallel simulation of relations would be facilitated when the simulations are dynamically aligned. Experiments 3 and 4 confirmed this prediction. In Experiment 3, participants responded faster when their limbs were manipulated symmetrically. In Experiment 4, people compared the outcomes of two mental rotations tasks faster when the absolute different between the two angles of rotation was smaller. Both results are consistent with the idea that dynamical alignment facilitates relation comparison. In Experiment 3, a symmetric manipulation of participants' hands meant that it was equally easy or hard to conduct the two perceptual-motor simulations and hence it was possible to align their start and end points. In Experiment 4, similar rotation angles meant that the durations of the two mental simulations would also be similar and it would be possible to dynamically align them. In this way, evidence from two different tasks, using different stimuli, was found in support of a prediction of our model of grounding relations in simulated actions.

Critics of embodiment

The embodied cognition approach promised to solve some of the fundamental problems of cognitive science, but not all cognitive researchers readily accepted it. The opponents of embodiment outline several main problems with it, some of which are also valid for our proposal that relations are embodied.

There is more to cognition than stimulus-response compatibility. A central claim of the embodiment approach is that cognition is grounded in the perceptual and motor systems. Typical experimental results which are interpreted as evidence that cognition is embodied demonstrate that irrelevant motor programs are activated during conceptual processing. An example of such a result is the finding that the recognition of objects with handles activates motor systems for manipulating them. The problem with such experiments is that they usually involve either neuro-imaging or behavioural experimentation based on stimulus-response compatibility. Mahon & Caramazza (2008) argued that neither kind of experiments can provide cogent evidence that cognition is embodied. The fact that a motor area of the brain is activated during conceptual processing does not entail that the motor programs are actually involved in this processing. A simpler 'disembodied' explanation exists - that such phenomena are due to spreading of activation between the conceptual, sensory and motor systems. The results of stimulus-response behavioural experiments can be explained in the same way – the activation of irrelevant motor systems interferes with the execution of the selected response action but does not actually play any role in conceptual processing. This criticism puts under question a great deal of the existing evidence in support of the embodiment view of cognition. However it is not relevant to any of the experiments described in this thesis. We have deliberately decided to use reaction time of verbal 'yes' responses as dependent measure in order to avoid spreading-activation explanations of our results. It is highly unlikely that regions of the brain involved in object manipulation and the production of 'yes' responses are such systematically connected that a stimulus-response compatibility conflict could arise. Therefore the effects found in Experiment 1, 2 and 3 can not be accounted for by simple associative mechanism. The result of Experiment 4 also can not be doubted on this ground because the motor programs potentially involved in mental rotation were the same in all experimental conditions and had no relation to the response action.

Embodiment is limited to concrete concepts. Theories of embodiment are often criticized that they can not account for abstract concepts and abstract reasoning (e.g., Clark, 1999; Dove, 2009). We have shown in this thesis that a cognitive function which is traditionally though to be highly abstract in nature – relation comparison – is constrained by the human body. However it remains an open question whether the same holds for more complex forms of relations and relational reasoning. For example, a particularly hard challenge to the embodied view of relations will be to show that higher-order relations are also embodied. One way to attack the problem is to exploit the idea of grounding abstract concepts in concrete experience by establishing metaphorical relations (Lakoff & Johnson, 1980). Our approach of modelling the embodiment of relations within a general analogy-making architecture (AMBR) makes the combination of embodied representations and metaphorical reasoning possible.

The neural mechanisms underlying embodied cognition are obscure. The theories of embodiment cognition make a number of predictions about the neural organization of cognition. However it is still debatable whether these predictions are supported by neurological findings. As we already discussed, the evidence from brain-imaging studies does not conclusively support the embodiment view as it is based on correlational data and bears alternative interpretations (Chatterjee, 2005). The neuropsychological evidence is controversial (Mahon & Caramazza, 2008; 2009).

The problem of explicating the neural mechanisms underlying embodied cognition is also valid for our proposal that relations are grounded in action. A further development of the embodied view of relations must specify how exactly embodied relational representations and relation reasoning are implemented in the brain.

Short-comings and limitations

Learning relations. A serious short-coming of the present work 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. The dynamics of the development of relational knowledge would generate additional predictions of the embodied view of relations and would highlight more ways in which sensorimotor experience and bodily constraints influence relational reasoning.

Using embodied relations in full blown analogy-making. One of the motivations to use

AMBR to model the grounding of relations was that such an approach would allow using embodied relation represenations 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 high-order relational structures.

No evidence that sensorimotor experience constraints relational reasoning. The idea of grounding relations in perceptual-motor simulations entails that the efficiency of thinking about relations is dependent on the ability to simulate actions and predict how they transform the perceptual input. Although this prediction is central to our proposal, we did not provide evidence that supports it, other than referring to the studies of Dixon & Dohn (2003), Trudeau & Dixon (2007) and Sam & Goldstone (2009).

Necessity. 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 can not preclude such an interpretation of the experimental results as we only found effects on reaction times and not on accuracy. Other experimental designs are needed in order to show that action plays a critical role in perceiving and reasoning with relations.

Generality. 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'.

Directions for future studies

There are a number of ways to extend the current work. Further developments of the computational model include a deeper exploration of the proposed idea that temporal role-filler binding could be carried out by the dynamic of action execution; modeling the process of learning embodied relational representations and the embodiment of high-order relations by establishing metaphorical relations to sensorimotor experience; showing how embodied relations are used in complex reasoning tasks.

The embodiment view of relations would be considerately strengthened if evidence is found that relations other than functional or spatial ones are embodied. Another direction for empirical investigation is concerned with revealing the dynamics of perceptual-motor simulations involved in understanding and reasoning with relations. For example, eye-tracking studies could be designed to show that, when comprehending functional relations between objects, people simulate interactions between them (Coventry et al., 2009). The same problem could be addressed by using designs from the studies of language comprehension, which showed that people mentally represent the orientation of objects implied by a verbal description (Stanfield & Zwaan, 2001). It also important to find evidence that bodily and environmental constrains affect relational reasoning not only in terms of processing speed, but also in terms of accuracy.

An important contribution to the embodied view of relations would be made by studying the neural mechanisms underlying embodied relational representations and reasoning processes. Neuropsychological studies of people suffering from apraxia or TMS studies could be particularly useful in providing conclusive evidence that relations are grounded in action.

References

Adamo, M., & Ferber, S. (2008). A picture says more than a thousand words: Behavioural and ERP evidence for attentional enhancements due to action affordances. Neuropsychologia.

Almeida, J., Mahon, B. Z., Nakayama, K., & Caramazza, A. (2008). Unconscious processing dissociates along categorical lines. Proceedings of the National Academy of Sciences of the United States of America, 105(39), 15214-15218.

Almeida, J., Mahon, B. Z., & Caramazza, A. (in press). The role of the dorsal visual processing stream in tool identification.

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, Vol. 43 (pp. 43-92). San Diego: Academic Press.

Beauchamp, M. S., Lee, K. E., Haxby, J. V., & Martin, A. (2002). Parallel visual motion processing streams for manipulable objects and human movements. Neuron, 34(1), 149-159.

Beauchamp, M. S., & Martin, A. (2007). Grounding object concepts in perception and action: evidence from fmri studies of tools. Cortex, 43(3), 461-468.

Bedny, M., Caramazza, A., Grossman, E., Pascual-Leone, A., & Saxe, R. (2008). Concepts are more than percepts: the case of action verbs. The Journal of neuroscience : the official journal of the Society for Neuroscience, 28(44), 11347-11353.

Berkely, G. (1709). A Essay Towards a New Theory of Vision . Project Gutenberg, http://www.gutenberg.org/etext/4722.

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.

Borghi, A. M., Bonfiglioli, C., Lugli, L., Ricciardelli, P., Rubichi, S., & Nicoletti, R. (2007). Are visual stimuli sufficient to evoke motor information? : Studies with hand primes. Neuroscience Letters, 411(1).

Brooks, R. A. (1991). Intelligence without representation. *Artif. Intell.* , *47* (1-3), 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.

Buccino G., Binkofski F., Fink G.R., Fadiga L., Fogassi L., Gallese V., Seitz R.J., Zilles K., Rizzolatti G., Freund H.J. (2001) Action observation activates premotor and parietal areas in a somatotopic manner: an fMRI study. European Journal of Neuroscience, 13, 400-404.

Buccino, G., Sato, M., Cattaneo, L., Roda, F., & Riggio, L. (2009). Broken affordances, broken objects: A tms study. Neuropsychologia, 47(14), 3074-3078.

Bundesen, C., & Larsen, A. (1975). Visual transformation of size. Journal of experimental psychology. Human perception and performance, 1(3), 214-220.

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.

Chatterjee, A. (2005). A madness to the methods in cognitive neuroscience? Journal of cognitive neuroscience , 17 (6), 847-849.

Clark, A. (1999). An embodied cognitive science? Trends in cognitive sciences , 3 (9), 345-351.

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.

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

Cooper, L. (1975). Mental rotation of random two-dimensional shapes. Cognitive Psychology, 7(1), 20-43.

Coventry, K. R. & Coventry, GarrodGarrod, S. C. (2004). Saying, Seeing and Acting. The psychological semantics of spatial prepositions. Taylor Francis, Psychology Press. Hove: New York

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.

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

Dietrich, E., & Markman, A. B. (2003). Discrete thoughts: Why cognition must use discrete representations. Mind & Language, 18(1), 95-119.

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).

Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. Cognition, 110(3), 412-431.

Ellis, R., Tucker, M., Symes, E., & Vainio, L. (2007). Does selecting one visual object from several require inhibition of the actions associated with nonselected objects? Journal of Experimental Psychology: Human Perception and Performance, 33(3), 670-691.

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.

Friedman, R., & Förster, J. (2002). The influence of approach and avoidance motor actions on creative cognition. Journal of Experimental Social Psychology, 38, 41-55.

Flusberg, S., Jenkins, G., & Boroditsky, L. (2009). Motor affordances in mental rotation: When minds reflect the world and when they go beyond. In N. Taatgen, & H. van Rijn (Eds.) Proceedings of the 31th Annual Conference of the Cognitive Science Society , (pp. 1453-1458). Austin, TX: Cognitive Science Society.

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

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

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.

Glenberg, A. M., & Kaschak, M. P. (2002). Grounding language in action. Psychonomic Bulletin and Review, 9, 558-565.

Gordon, P. C., & Moser, S. (2007). Insight into analogies: Evidence from eye movements. Visual Cognition, 15(1), 20-35.

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

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. Trends in neurosciences , 15 (1), 20-25.

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.

Hickok, G. (2009). Eight problems for the mirror neuron theory of action understanding in monkeys and humans. Journal of cognitive neuroscience, 21(7), 1229-1243.

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

Hommel, B., Müsseler J, Aschersleben G, & Prinz W. (2001) The theory of event coding (TEC): A framework for perception and action planning. Behavioral and Brain Sciences, 24, 849-878.

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. (1999). Binding problem. In R.A. Wilson and F.C. Keil (Eds.), The MIT Encyclopedia of the Cognitive Sciences, MIT Press.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. Psychological Review, 99(3), 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.

Hurley, S. (2008). The shared circuits model (scm): How control, mirroring, and simulation can enable imitation, deliberation, and mindreading. Behavioral and Brain Sciences, 31(01), 1-22.

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.

Kaschak, M. P., Madden, C. J., Therriault, D. J., Yaxley, R. H., Aveyard, M. E., Blanchard, A. A., & Zwaan, R. A. (2005). Perception of motion affects language processing. Cognition, 94, B79-B89.

Kaup, B., Yaxley, R. H., Madden, C. J., Zwaan, R. A., & Ludtke, J. (2007). Experiential simulations of negated text information. Quarterly journal of experimental psychology (2006), 60(7), 976-990.

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.

Kokinov, B., Vankov, I., Bliznashki, S. (2009). How analogy could force re-representation of the target and inhibition of the alternative interpretation. In: Kokinov, B., Holyoak, K., Gentner, D. (eds.). New Frontiers in Analogy Research. Sofia: NBU Press.

Kosslyn, S. (1994). Image and Brain: The Resolution of the Imagery Debate. MIT Press, Cambridge, MA.

Kroger, J. K., Holyoak, K. J., & Hummel, J. E. (2004). Varieties of sameness: the impact of relational complexity on perceptual comparisons. Cognitive Science: A Multidisciplinary Journal, 28(3), 335-358.

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.

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.

Lingnau, A., Gesierich, B., & Caramazza, A. (2009). Asymmetric FMRI adaptation reveals no evidence for mirror neurons in humans. Proceedings of the National Academy of Sciences of the United States of America, 106(24), 9925-9930.

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 psychology. 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.

Mahon, B., & Caramazza, A. (2009). Concepts and categories: a cognitive neuropsychological perspective. Annual review of psychology , 60 (1), 27-51.

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

Moller, R., & Schenck, W. (2008). Bootstrapping cognition from behavior-a computerized thought experiment. Cognitive Science: A Multidisciplinary Journal, 32(3), 504-542.

Myung, J.-Y. Y., Blumstein, S. E., & Sedivy, J. C. (2006). Playing on the typewriter, typing on the piano: manipulation knowledge of objects. Cognition , 98 (3), 223-243.

Negri, Gioia, Rumiati, Raffaella, Zadini, Antonietta, Ukmar, Maja, Mahon, Bradford, Caramazza, & Alfonso (2007). What is the role of motor simulation in action and object recognition? evidence from apraxia. Cognitive Neuropsychology, 24(8), 795-816.

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.

Pecher, D., Zeelenberg, R., & Barsalou, L.W. (2003). Verifying properties from different modalities for concepts produces switching costs. Psychological Science, 14, 119-124.

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.

Pulvermüller, F., M. Haerle, & F. Hummel (2001). Walking or Talking?: Behavioral and Neurophysiological Correlates of Action Verb Processing. Brain and Language, 78, 143–168.

Regier, T. (1996). The human semantic potential: spatial language and constrained connectionism. Cambridge, MA: MIT Press.

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.

Richardson, D., & Matlock, T. (2007). The integration of figurative language and static depictions: an eye movement study of fictive motion. Cognition, 102(1), 129-138.

Richardson, D. C., Spivey, M. J., Barsalou, L. W., & McRae, K. (2003). Spatial representations activated during real-time comprehension of verbs. Cognitive Science: A Multidisciplinary Journal, 27(5), 767-780.

Rizzolatti, G. & Cragehero, L. (2004) The mirror-neuron system. Annual Review in Neuroscience, 27, 169-192.

Rizzolatti G., Fogassi, L., & Gallese, V. (2001). Neurophysiological mechanisms underlying the understanding and imitation of action. Nature Reviews Neuroscience, 2, 661-670.

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.

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). E-Prime user's guide. Pittsburgh: Psychology Software Tools Inc.

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

Simon, J., & Rudell, A. (1967). Auditory S–R compatibility: The effect of an irrelevant cue on information processing. Journal of Applied Psychology, 51, 300-304.

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

Siskind, J. M. (1994). Grounding language in perception. Artificial Intelligence Review, 8(5), 371-391.

Siskind, J. (2003). Reconstructing force-dynamic models from videoA sequences. Artificial Intelligence, 151(1-2), 91-154.

Stoet, G. & Hommel, B. (1999). Action planning and the temporal binding of response codes.. Journal of Experimental Psychology: Human Perception and Performance, 25, 1625-1640.

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.

Stanfield, R. A., & Zwaan, R. A. (2001). The effect of implied orientation derived from verbal context on picture recognition. Psychological Science, 12(2), 153-156.

Stanford Encyclopedia of Philosophy. Online resourse. http://plato.stanford.edu/entries/mental-imagery.

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.

Symes, E., Ellis, R., & Tucker, M. (2005). Dissociating object-based and space-based affordances. Visual Cognition, 12(7), 1337-1361.

Symes, E., Ellis, R., & Tucker, M. (2007). Visual object affordances: object orientation. Acta psychologica, 124(2), 238-255.

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. (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 psychology. Human perception and performance, 24(3), 830-846.

Turella, L., Pierno, A. C., Tubaldi, F., & Castiello, U. (2009). Mirror neurons in humans: consisting or confounding evidence? Brain and language, 108(1), 10-21.

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.

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 Series in Cognitive Science.

Vankov, I., & Kokinov, B. (2010) The role of action in perceiving and comparing functional relations. Proceedings of the 32nd Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society.

Vingerhoets, G., Vandamme, K., & Vercammen, A. (2009). Conceptual and physical object qualities contribute differently to motor affordances. Brain and Cognition, 69(3).

Wexler, M. (1998). Motor processes in mental rotation. Cognition, 68(1), 77-94.

Willems, R. M., Hagoort, P., & Casasanto, D. (2009). Body-specific representations of action verbs: Neural evidence from right- and left-handers. NeuroImage, 47

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.

Yaxley, R. H., & Zwaan, R. A. (2007). Simulating visibility during language comprehension. Cognition, 105(1), 229-236.

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

APPENDIX A. Stimuli used in Experiments 1, 2 & 3

Target trials

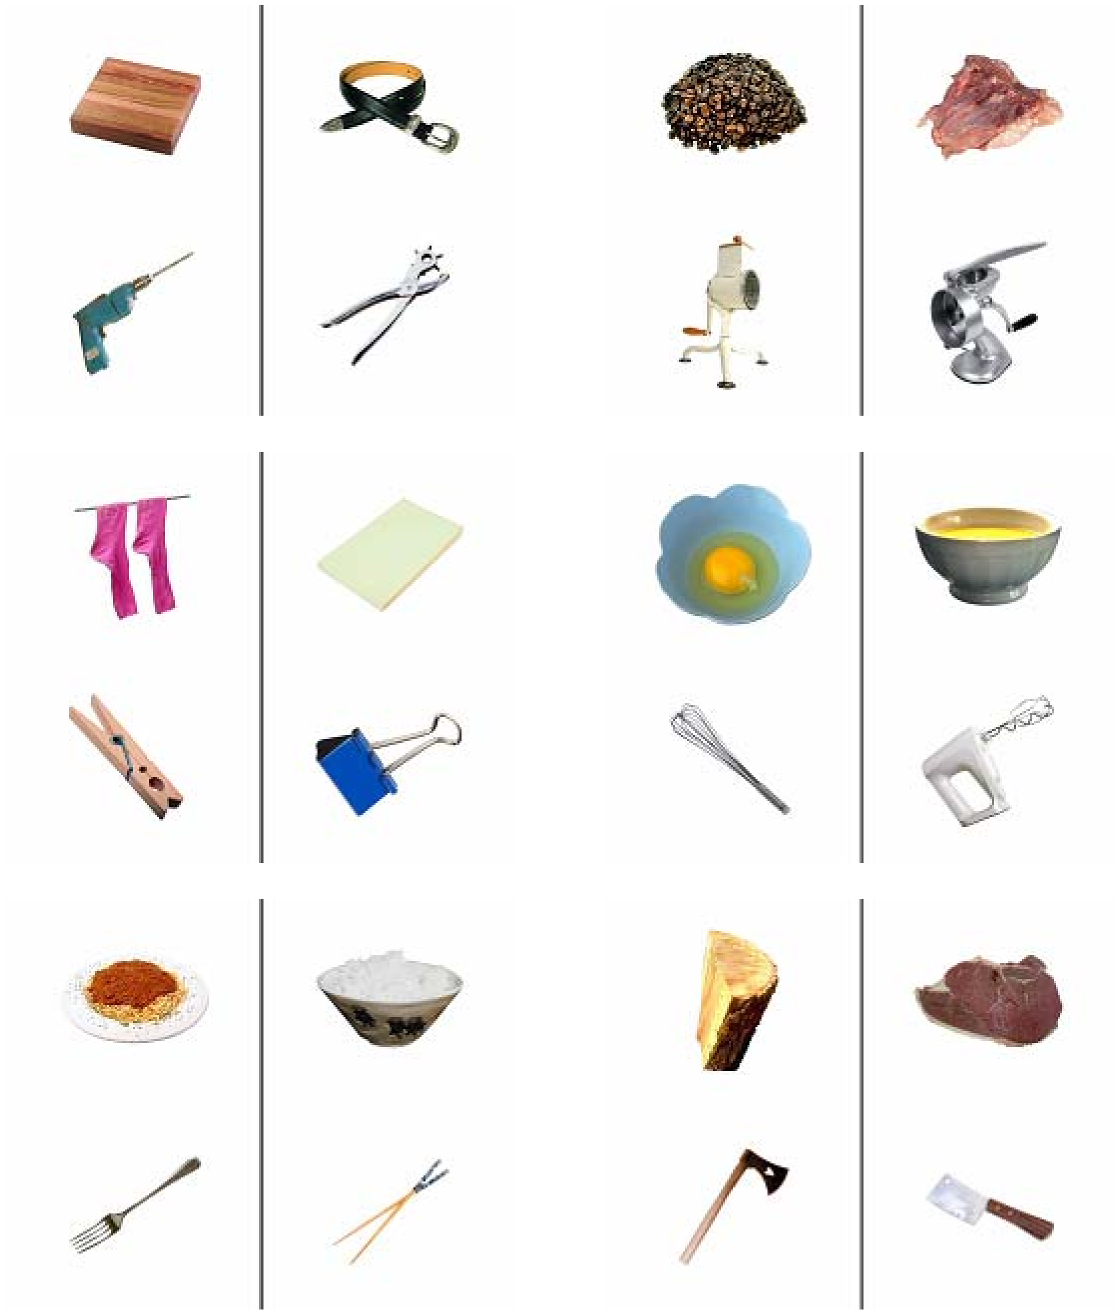







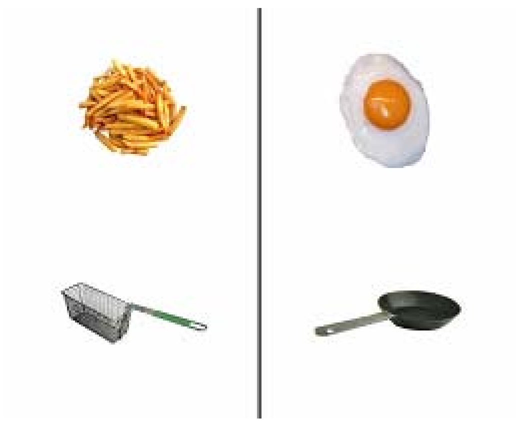


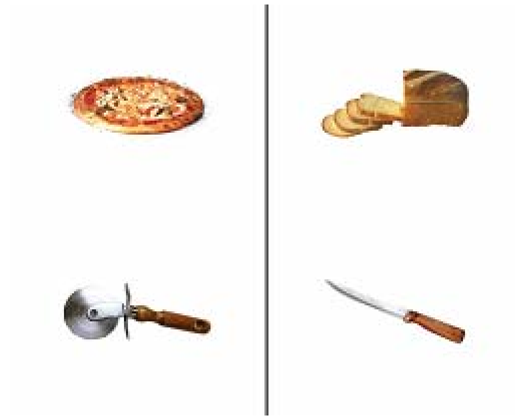

Fillers


БЪЛГАРИЯ
BULGARIA
















APPENDIX B. Stimuli used in Experiment 4