Distribution and Association: Modeling Two Fundamental Principles in Cognitive Control

Holger Schultheis (schulth@informatik.uni-bremen.de)
University of Bremen, Faculty 03: Mathematics/Computer Science
P.O. Box 330 440, 28334 Bremen

Abstract

Cognitive control is pervasive in human behavior. No matter what task one is performing, control of the sequence of one's actions seems indispensable to accomplish the task. Therefore, control is one of the key issues in understanding human behavior. Moreover, knowledge about how control is exerted in humans may potentially help in designing improved artificial systems, since the control performance humans exhibit is currently unrivaled by artificial systems. Both understanding human control behavior and building improved artificial systems could be achieved by a sufficiently precise computational model of control which, however, currently does not exist. This paper devises a new computational model. In doing so, two fundamental principles of cognitive control in humans are identified. Based on these insights a new model of control which implements those principles is developed.

Introduction

In everyday behavior[1] for every action one takes there are infinitely many other actions possible which one ignores. Instead of reading this text, for instance, you could read something else, contemplate your own research ideas or do anything else you can imagine.

Consequently, the question arises "What makes one take the action one takes and not one of the other possible actions?" or, in other words, "How is one's behavior controlled?". To find an answer to this question is of major importance to psychology, since the issue of control is at the heart of not only some special tasks, but of all human behavior. Moreover, answering this question is also of interest to the field of artificial intelligence, because the control performance of humans exhibits an efficient combination of goal directedness and simultaneous flexibility unrivaled by current artificial cognitive agents/systems. Thus, understanding how human behavior is controlled may help in building better artificial systems. Accordingly, there has been considerable effort to both model control (e.g., Kieras, Meyer, Ballas, & Lauber, 2000; Kimberg & Farah, 2000; Meiran, 2000) and to answer the question of how control is realized in humans (e.g., Allport, Styles, & Hsieh, 1994; Baddeley, 2002; Gopher, 1996; Goschke, 2004; Kluwe, 2000; Norman & Shallice, 1986).

An ad hoc answer would be that one's will causes the action to happen. On a closer look, however, there are several problems with this account of control. First of all, assuming that will controls behavior raises the question who or what controls will. Thus, introducing something like will only shifts the problem, leading, in the end, to an infinite regress. A second reason to question that control is exerted solely by will is the fact that in certain situations behavior is seemingly controlled by other factors. For instance, in one classical experimental paradigm, the antisaccade task, participants have to fixate the middle of a computer screen until either on the left or right part of the monitor some stimulus (e.g., a cross) appears. Participants are instructed to look, as soon as the stimulus is presented, at the part of the monitor where the stimulus is not shown, i.e., away from the stimulus. As, among others, Kane, Bleckley, Conway, and Engle (2001) have shown, participants despite their will and effort to look away from the stimulus quite often look towards it. This clearly indicates that there have to be additional determinants of control besides will.

As a result, during the last century researchers have tried to clarify the functioning of control of behavior and to explain it without referring to some homunculus. These efforts notwithstanding, Hommel, Daum, and Kluwe (2004) point out that most of how control is exerted is still unknown. This is, as Monsell and Driver (2000) remark, in part due to the lack of a sufficient computational model of control processes which could potentially both constitute a theoretical framework to integrate experimental results and provide a possibility to test different theoretical accounts. Since, as already mentioned, such a model would also be of advantage to the design of artificial cognitive systems its development seems a worthwhile endeavor.

Consequently, the conceptualization and implementation of such a model is the aim of the research presented in this paper. To this end, psychological work concerning control is shortly summarized and two fundamental results emerging from this work are identified in the following section. Subsequently, previous models of control are critically discussed with respect to the two fundamental results before a new modeling approach is proposed. As a final point, open issues regarding the new approach are discussed.

Fundamental Results From Psychology

Despite its importance, control has become a major research topic only comparatively recently, i.e., in the late 1980ies and early 1990ies. Therefore, experimental results regarding control in general which are detailed enough to be used for modeling purposes are virtually nonexistent. Instead, aspects of control have been examined in more depth almost merely with respect to particular activities as, for instance,

[1] Here and in the following the terms behavior and action will be used in their broadest sense, i.e., including also purely mental behavior/action.

task switching see Monsell (2003, for a review) or the antisaccade task (cf. Kane et al., 2001). Nonetheless, results emerging from this research point to fundamental principles of how control is realized in humans. Two of these fundamental principles, namely distribution and association, will be described in the following subsections.

Control is Distributed

It has been an ongoing debate whether control is realized by a unitary component or whether it is the result of many different (probably interacting) components. In neuropsychology, for instance, it has been argued (cf. Norman & Shallice, 1986) that all control is realized by one single brain structure, namely Prefrontal Cortex. Likewise, there are classical psychologists, as Gopher (1996), advocating a single control component.

Contradicting to such a unitary component assumption are results from taskswitching experiments (see, e.g., Allport et al., 1994; Monsell, 2003): Certain experimental manipulations do influence task switching performance in a way inconsistent with a unitary component of control. Moreover, evidence from neuroimaging studies, as the one conducted by Garavan, Ross, Li, and Stein (2000), suggest that the Prefrontal Cortex is not the only brain structure involved in control, but that there are many different brain regions distributed over the cortex. Consequently, a majority of researchers (e.g., Allport et al., 1994; Engel, Bertel, & Barkowsky, 2005; Goschke, 2004; Kluwe, 2000) argue that control is the result of the interaction of several different components or, more concisely, that control is distributed.

Control is Associative

The importance of associations to the problem of control stems from the fact that objects and actions or actions and actions may be associated with each other. Accordingly, stimuli in the environment or the mind may elicit certain actions, i.e., stimuli may induce behavior. This is especially evident every time stimulustriggered actions are inappropriate. One particular example would be looking towards the stimulus in the antisaccade task described above.

Although we may be only aware of the influence of stimuli when this leads to error, recent research, as, e.g., Hommel et al. (2004) remark, has shown that associative control is quite pervasive. Unfortunately, however, it is not easy to conceive whether and how association alone enables reasonable behavior in all situations. Considering the large amount of entities that may be associated with one stimulus it is not obvious how such a system of associations could be configured such that it responds correctly in every situation in which this stimulus occurs. In accord with these considerations, psychological work (see Goschke, 2004; Hommel, Ridderinkhof, & Theeuwes, 2002; Kieras et al., 2000) indicates that association, although an important factor, is not the only determinant of control. Instead, a generally accepted account which can be traced back to the beginning of the last century (Ach, 1905) assumes that higher level control processes like intention modulate associative control. If, for example, one's intention is to determine the sum of two digits this intention will lead to a preselection of associations relating digits with the appropriate action (namely adding). Given this preselection the occurrence of two digits will—associatively, and due to the preselection almost exclusively—elicit the action of adding the two digits.

Such a theory can still account for the flexibility of human behavior which can quickly adapt to new situations, while at the same time being able to explain the goaldirectedness, commonly observed in human behavior. Yet, the introduction of the additional factor (i.e., intention) destroys the simplicity of the associative approach, since this factor is potentially homunculus like and needs further explanation. Although the latter issue is not completely settled yet, one idea emerging throughout the field is that intentions stem from perceptions (see, e.g., Kluwe, 2000).

In summary, there is considerable evidence showing that control is achieved by an interaction of bottomup (associative, stimulusdriven) and topdown (intention) processes lending the flexibility and directedness to human behavior commonly observed.

Having identified two fundamental principles, namely distributedness and associativeness, acting on the control of human behavior, in the following two sections the impact of these fundamentals on modeling control is considered. To this end, existing conceptions of control will be evaluated regarding the principles before a new approach is proposed.

Previous Conceptions of Control

Since control is at the heart of all human behavior, every model of behavior has to (explicitly or implicitly) implement control in some way or the other. In this section the most renowned and successful of these modeling approaches will be examined with respect to their implementation of control. More precisely, the models considered in this section comprise EPIC (Kieras & Meyer, 1997), ACTR (Anderson et al., 2004), and Soar (Newell, 1990) which will be discussed in turn.

EPIC

The EPIC architecture is composed of several modules, called processors, conceived as running in parallel and cognition is realized as a sequence of executed productions. Only one of the modules, namely the cognitive processor, contains and executes these production, i.e., only one processor implements cognition proper. In this framework control is achieved by a special subset of productions not being concerned with realizing some task, but scheduling task execution similar to how an operating system schedules tasks on a computer (cf. Kieras et al., 2000). Such a control conception seems unsatisfactorily with respect to the above identified principles of control. First of all, control is not distributed in EPIC, since there is only a single module implementing control, namely the cognitive processor. Second, since control is explicitly exerted by a set of productions, there is no associative but a deterministic preselection of stimulusresponse dispositions (i.e., productions). Deterministic preselection, however, is unable to model, in a general way, behavioral lapses like looking towards the stimulus in the antisaccade task (see above). It is, of course, possible to model such erroneous behavior explicitly in EPIC. Yet, the need for explicit modeling shows that the model is unable to explain control errors with its general control mechanisms. In summary, EPIC fails to implement both distribution and association in a satisfactory manner.

ACT-R

This architecture shows several similarities to EPIC. First, the information processing system is conceptualized as being composed of modules[2]. Second, information processing is modeled by firing productions. Finally, there is a single module firing productions. In contrast to EPIC, however, rather than assuming a special set of productions doing the controlling, control is realized implicitly by the same productions that are implementing the task proper. This is achieved by positing a special module, namely the—unfortunately not specified in detail—intention module, which generates and influences the current goal. The current goal, again, plays an important role in selecting productions. Each production is tightly bound to a goal and can only be fired if its goal agrees with the one currently pursued. Thus, depending on the current goal only a subset of all available productions can be triggered by additional stimuli (e.g., information retrieved from memory). That is, the goal in ACT-R realizes some preselection of the kind which can be found in human control behavior. Despite this similarity preselection in the model differs in a substantial aspect from preselection in humans: Selection in the model is deterministic and not associative. Accordingly, ACT-R suffers from the same problems the determinism of production selection causes in EPIC. As, in addition, control in ACT-R is mainly conceptualized as not distributed[3], the architecture does not seem to be sufficiently in accord with knowledge about human control processes.

Soar

Despite many differences between Soar and the two afore mentioned architectures it bears considerable resemblance to them regarding control. Not only that cognition proper is realized as production firing, but also the execution of productions is carried out by a single component of the system. Furthermore, as in ACT-R, at any point in processing only a subset of all possible productions is suitable for application (i.e., Soar employs some kind of preselection). More precisely, every production belongs to a so called problem space and can only be fired if this problem space is identical to the one currently active. Accordingly, preselection is deterministic in Soar, too. Thus, Soar does not seem to put into operation the fundamental principles of human control behavior any better than EPIC or ACT-R.

Consequently, notwithstanding their unarguable value and success in modeling human behavior, none of the three considered architectures appears to be a sufficiently accurate model of control processes, because they all fail to satisfactorily implement the fundamental principles of associativeness and distributedness observed in the control of human behavior. Due to the importance of an accurate control model to both the field of psychology and artificial intelligence it therefore seems desirable to improve on the existing models or develop a new one to obtain models more in accord with the fundamental psychological results. In the following section such an improved model is proposed before some concluding remarks provide insight on open research questions and future work.

An Improved Model

Due to the ability of the above architectures to model a wide range of human behavior it seems to be a sensible strategy in designing a new model to keep as much of the previous approaches as possible and change only those aspects necessary to improve model adequacy regarding control. To this end, the proposed model will adopt two principles common to all of the above reviewed architectures, namely modularity and use of productions. More precisely, in the new approach it is assumed that (a) the cognitive system is comprised of individual parts (i.e., modules) each of which serves some distinct aspect of cognition and (b) cognitive processes in general can be modeled by productions.

In contrast to the previous approaches, however, every module has its own set of productions specific in their working characteristics to the respective module. That is, instead of having one single coordinating component gathering and integrating the information from all other modules every module is able—via the productions—to gather, integrate and make use of the information from all other modules. In particular, there is no single instance coordinating the interaction between the different modules: Every module acts solely on the basis of its applicable productions and not according to what some central coordinator instructs. This conception of the modules entails that other than in EPIC and keeping with ACT-R and Soar there is no set of special productions coordinating the overall information processing. Rather each module focuses on its own field of competence—delineated by its set of productions—with the overall coordination emerging from the interaction of all (locally focused) modules.

By such conceptualizing the modules the first of the detected fundamental principles of control in human behavior, namely distribution of control, is realized by the new architecture. Yet, distributing control, though a first step towards a more accurate model of human control behavior, is in itself not sufficient. To begin with, given this distributed, heterarchical conception of coordination the question arises how reasonable overall performance of the model can be achieved, i.e., how exactly coordination is achieved without falling back on a single coordinating module. Furthermore, the second fundamental principle identified above, namely the associativeness of control, still needs to be taken into account in the model. So it seems that there are still two problems to be solved. Yet, as detailed in the next paragraphs, it turns out, that by introducing associativeness in an appropriate way the problem of coordination is solved at the same time. In other words, the coordination problem resulting from distributed control is remedied by implementing the second fundamental principle.

A major prerequisite for introducing association principles into the model is to change the way productions are selected. More precisely, the influence current context information exerts on production selection has to be strengthened. In EPIC, ACT-R, and Soar, the context information (e.g., the current goal and/or objects in working memory) is only important for deciding whether a production can be applied or not. If the

[2]Actually, several modules used in EPIC have been adopted for ACT-R.

[3]Although Anderson et al. (2004) admit that goals may be represented in a distributed way, (a) they currently are not and (b) production firing is realized by a single component.

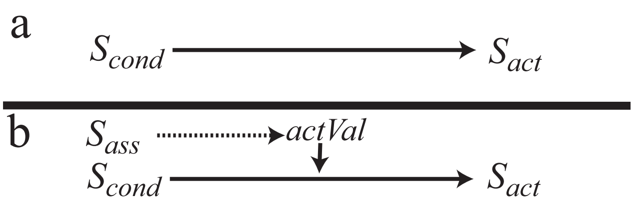

Figure 1: Productions as commonly used (a) and as newly proposed (b). Dashed arrows indicate associative influence.

left-hand side of a production matches the context information it can be applied.

Stated more formally, this way of production selection can be characterized as follows:

- Context information comprises several discrete blocks of information, called *elements*, henceforth denoted by e_i. Furthermore, the elements currently constituting the context information are a subset of the set of all possible elements E: context information $\subseteq E = \{e_1, \ldots, e_i, \ldots\}$.
- The left-hand side of a production tests for the existence of particular elements in the context. Therefore, the left-hand side of a production can be conceived as a set of elements, henceforth denoted by S_{cond}, i.e., $S_{cond} \subseteq E$.
- The production is assumed applicable only if $S_{cond} \subseteq$ context information.
- In case a production is fired, its right-hand side, that is, a set of actions is executed. For every production this set is a subset of all possible actions A. Denoting single actions by a_i and the right-hand side of productions by S_{act} one gets: $S_{act} \subseteq A = \{a_1, \ldots, a_i, \ldots\}$.

Using these formalisms, productions as used by EPIC, ACT-R, and Soar can be represented as shown in Figure 1a.

The third of the above formulas implies that in these three architechtures the influence of context information on production selection is strictly deterministic: Either the production is applicable or not. Since, as the analysis in the preceeding section has shown, such a deterministic selection of productions is not in accord with the associativeness observed in the control of human behavior, the determinism needs to be weakened or enhanced by an additional mechanism.

The respective mechanism proposed in the new model comprises three changes to the usual production selection. First, with each production an additional value, denoted by $actVal$, indicating how strongly the current context suggests the application of this production is stored. In case there is more than one production applicable (in the deterministic sense) the one with the highest activation value is fired. In particular, this production activation value is assumed to be governed by the same laws as the activation values of entities in other associative systems (e.g., memory), i.e., (a) use of a production increases its activation, (b) activation decays with time, (c) this decay is exponentially decelerated leading to an asymptotic minimum of activation, called baseline activation and (d) activation does randomly vary over time (noise). Second, apart from S_{cond} and S_{act} every production receives a third set S_{ass} containing all elements with which this production is associated: $S_{ass} \subseteq E$. Third, the presence of one of the elements on this third list of a production in the context increases the activation of that production. The type of production resulting from these three changes can be represented as depicted in Figure 1b. Given such productions, context information not only determines the (all or none) application of productions but also influences the activation value of productions and thereby their chance of being fired. Hence, the first aspect of the associativeness of control, namely bottom-up influence, is implemented in the proposed model. As already mentioned above, however, besides bottom-up processes, top-down processes have a considerable associative bearing on the organization of behavior. Accordingly, to build an accurate control model, intentions/goals and their impact on behavior control have to be considered as well. To this end, it is assumed that one of the elements which are part of the context is representing the current intention. Such an element is called goal and denoted by g_i. More precisely, it is assumed that there is a distinct subset G of E comprising all elements which are goals. As all other context elements, but to a greater extent, the goal, too, activates those productions which are associated with it. What is more, the influence of the goal is conceptualized as being solely associative. That is, rather than strictly excluding productions which are not associated with the current goal the chance of firing those productions which are associated with the goal is increased. By this means the experimentally observed fact that some actions are executed despite and not in accord with the current goal can be explained: Due to context information some production not associated with the goal is more strongly activated than any production which is associated with the goal. This conceptualization implies that goals are only allowed to appear in S_{ass} but not in the left-hand side of a production. Formally, this amounts to a redefinition of S_{cond} for the new productions, i.e., $S_{cond} \subseteq E \backslash G$.

To complete the overview of the proposed model it remains to further detail that part of the model which up until now has been termed "context", since this context contains all elements activating and selecting productions and is of major importance for the realization of the associative principles. Technically speaking the context is a bus, i.e., a system of connections between the different modules[4]. In particular, there is only one bus/context to which all of the modules have full access (see Figure 2). Any module can therefore put any results of its internal workings, denoted by M_{out}, onto the bus and all elements currently on the bus potentially influence the production selection in every module. As soon as a new element is put onto the bus it gets some initial activation which will decay over time. If the activation of a context element falls below a certain threshold it will—assuming that it is not accessible for the modules any more—be removed from the bus. In starting with an initially (high) activation value which reduces over time it is guaranteed that the most recent information has the greatest influence on behavior. Thus, the context enables the distinct modules to exchange information and thereby, due to the associative principles, at the same time influence each others' processing.

[4]Due to its functionality the bus as used in the proposed model is quite similar to a blackboard.

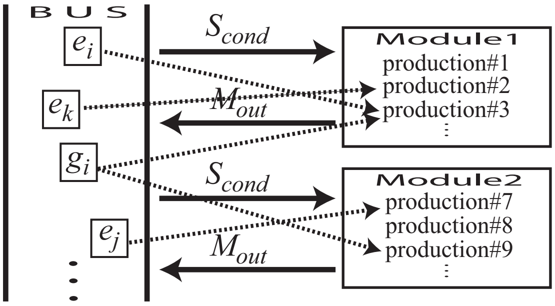

Figure 2: Structure of the proposed model. Dashed arrows indicate associative influence.

Modeling Antisaccade Behavior

To give an impression of the capabilities of the proposed architecture to account for empirical data, in the following, a model run for an experimental task will be presented. More precisely, a model of human behavior in the antisaccade task as employed by Kane et al. (2001, experiment 2) and the models sequence of actions while working on one trial of that task will be described. In particular, it will be shown that the model accounts nicely for the fact that subjects in about one third of all trials, despite their goal to look away from a cue shortly preceding the target, look towards the cue.

A trial of the antisaccade task starts with a blank screen. Then, preceded by a fixation screen and again a blank screen, the cue is presented left or right to the center of the screen. This cue reliably predicts the region where the target—which can be one of three letters—will appear later on: If the cue is to the left the target will be to the right and vice versa. With disappearance of the cue a blank screen followed by the target is presented. Subsequently the target is masked until a key is pressed by the subject to indicate which letter has been recognized. Regarding control, the crucial aspect of this task is in which direction the first saccade after appearance of the cue occurs. Accordingly, in modeling behavior in the antisaccade task, this aspect of the task has been considered in more detail than aspects less important to control (as, e.g., how identification of the target letter is realized).

The resulting model has the following properties: (1) It comprises three modules, namely working memory, a vision module, and a motor module; (2) every element entering the bus has activation 1; (3) element activation decays according to $\exp(-5t)$, where t is time in seconds; (4) if the activation of an element falls below 0.5 it is removed from the bus; (5) firing any production takes 50 ms; (6) baseline activation of highly practiced and unpracticed productions is 3 and 1, respectively; and (7) production activation noise is distributed normally with mean 0 and variance 0.5. Working memory contains one production to write the goal (do the antisaccade task) to the bus and one to remove the goal when the task is finished. Both are associatively linked to the goal, i.e., their S_{ass} contain the goal. The motor module houses just one production which initiates a key press corresponding to the recognized target and is also associated with the antisaccade goal. The vision module, finally, comprises five productions. The first production perceives the objects currently displayed on the screen, i.e., given a particular screen writes the objects indicating their type (fixation, cue, target, or mask) and position (left, center, or right) to the bus. The second and third production initiate a saccade to the left or right, respectively, if element(s) on the bus indicate that the cue is to the left or right. Those productions, since they can be assumed to represent automatic behavior are not bound to a certain goal (cf. Kane et al., 2001) and thus their S_{ass} lists do not contain the antisaccade goal. In contrast, the fourth and fifth production initiate a saccade to the right or left, respectively, if element(s) on the bus indicate that the cue is to the left or right. As these productions are not automatic, but specific to the antisaccade task, they are associated with the goal.

Given these properties and productions, processing of one antisaccade trial by the model results in the following trace:

```
TIME  EVENTS AND BUS STATE
   0: BUS = {goal}
 400: Fixation screen appears; BUS = {goal}
 450: BUS = {goal, percept(fixation, center)}
 600: Blank screen; BUS = {goal, percept}
 650: Cue appears; BUS = {goal, percept}
 700: BUS = {goal, percept(fixation, center),
             percept(cue, left)}
 900: Blank screen; BUS = {goal}
 950: Target appears; BUS = {goal}
1000: BUS = {goal, percept(target, right,
             'B')}
1050: Mask appears; BUS = {goal, percept}
1100: BUS = {goal, percept(mask, right),
             percept(target, right, 'B')}
1450: Blank screen; BUS = {goal,
                    percept(mask, right)}.
```

As can be seen, the goal is always present and indeed actively refreshed by the production in working memory. The interplay of this refreshing and the decay of element activation results in a constant goal activation of approximately 2.4. Any time a stimulus is displayed on the screen (e.g., after 400 and 650 ms) the first vision production writes the perceived object to the bus (e.g., at 450 and 700 ms). A soon as the cue percept is available on the bus (at 700 ms) productions two (p_2) and four (p_4) of the vision module compete for saccading to the left or right, respectively. In deciding which production to fire, their activation values are computed as $p_2: 3 +$ noise and $p_4: 1 + 2.4 +$ noise, where the '2.4' in the second equation stems from the activation of the goal p_4 is associated with.

Repeatedly running this model on antisaccade trials reveals that p_2 will be selected in approximately 35% of all trials. This value agrees nicely with the behavior human subjects show in the antisaccade task (see Kane et al., 2001). Thus, in accord with the theoretical considerations, the new architecture is capable of accurately modeling cognitive control.

Conclusions and Future Work

With the aim of obtaining an accurate model of human control a thorough review of psychological literature regarding control has been carried out. In doing so, two fundamental principles of human control behavior, namely distributedness and associativeness, have been pinpointed. Since all of the previously proposed models fail to realize these principles a new model has been devised. This model, as detailed in the above section, more faithfully mirrors the mechanisms at work in humans and thus constitutes a more accurate model of control as the previous approaches.

However, even though the new model is more accurate than earlier ones, there are still some aspects of it where further improvement seems possible. For one, the way goals emerge during information processing in humans is not satisfactorily modeled. At the moment goals are assumed to be either set from the outside when starting the model or may be—once the model is running—updated by productions in the modules. Although productions, i.e., cognition, may well be one factor determining goals it does not seem to be the only determinant (see, e.g., Norman & Shallice, 1986). Instead, motivational and perceptual influences are thought to be of major importance in setting up goals. Besides, the two principles of the control of human behavior identified and implemented in the new model are based to a great extent on research regarding particular activities like task switching (cf. Monsell, 2003) or the antisaccade task (cf. Kane et al., 2001). Accordingly, the accuracy of the new approach as a model of control can be assumed to be especially high when modeling those activities. It is less obvious, though, how good the model is able to account for human behavior in situations affording different control skills as, for instance, monitoring and correcting errors in one's own behavior.

To develop a control model even more accurate both the question of how well the new approach generalizes to additional control skills and the question of how to generate goals have to be taken into account. Accordingly, the next steps will be to further improve the model proposed above by testing its generalization capabilities and by extending it to include (a) mechanisms which in a reasonable way generate goals and (b) mechanisms—if need may be—allowing to model other control skills.

Acknowledgements

In this paper work done in the project R1-[ImageSpace] of the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition is presented. Funding by the Deutsche Forschungsgemeinschaft (DFG) is gratefully acknowledged. I also thank Thomas Barkowsky and Sven Bertel for helpful comments on earlier drafts of this document.

References

Ach, N. (1905). *Über die Willenstätigkeit und das Denken*. Göttingen: Vandenhoeck & Ruprecht.

Allport, A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set : Exploring the dynamic control of tasks. In C. Umilta & M. Moscovitch (Eds.), *Attention and Performance XV.* Cambridge, MA: MIT Press.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036 - 1060.

Baddeley, A. (2002). Fractionating the central executive. In D. T. Stuss & R. T. Knight (Eds.), *Principles of frontal lobe function.* New York: Oxford University Press.

Engel, D., Bertel, S., & Barkowsky, T. (2005). Spatial principles in control of focus in reasoning with mental representations, images, and diagrams. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, & T. Barkowsky (Eds.), *Spatial Cognition IV. Reasoning, Action, Interaction.* Berlin: Springer.

Garavan, H., Ross, T. J., Li, S.-J., & Stein, E. A. (2000). A parametric manipulation of central executive functioning. *Cerebral Cortex*, *10*, 585 - 592.

Gopher, D. (1996). Attention control: explorations of the work of an executive controller. *Cognitive Brain Research*, *5*(1 - 2), 23 - 38.

Goschke, T. (2004). Vom freien Willen zur Selbstdetermination. *Psychologische Rundschau*, *55*(4), 186 - 197.

Hommel, B., Daum, I., & Kluwe, R. H. (2004). Exorcizing the homunculus, phase two: Editors' introduction. *Acta Psychologica*, *115*, 99 - 104.

Hommel, B., Ridderinkhof, K. R., & Theeuwes, J. (2002). Cognitive control of attention and action: Issues and trends. *Psychological Research*, *66*, 215 - 219.

Kane, M. J., Bleckley, M. K., Conway, A. R. A., & Engle, R. W. (2001). A controlled-attention view of working-memory capacity. *Journal of Experimental Psychology: General*, *130*(2), 169 - 183.

Kieras, D. E., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, *12*, 391 - 438.

Kieras, D. E., Meyer, D. E., Ballas, J. A., & Lauber, E. J. (2000). Modern computational perspectives on executive mental processes and cognitive control: Where to from here? In S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII: Control of cognitive processes.* Cambridge, MA: MIT Press.

Kimberg, D. Y., & Farah, M. J. (2000). Is there an inhibitory module in the prefrontal cortex? Working memory and the mechanisms underlying cognitive control. In S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII: Control of cognitive processes.* Cambridge, MA: MIT-Press.

Kluwe, R. H. (2000). Steuerung des Denkens und Handelns. *Zeitschrift für Psychologie*, *208*(1 - 2), 1 - 31.

Meiran, N. (2000). Modeling cognitive control in task switching. *Psychological Research*, *63*, 234 - 249.

Monsell, S. (2003). Task switching. *TRENDS in Cognitive Sciences*, *7*(3), 134 - 140.

Monsell, S., & Driver, J. (2000). Banishing the control homunculus. In S. Monsell & J. Driver (Eds.), *Attention and Performance XVIII: Control of cognitive processes.* Cambridge, MA: MIT Press.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Norman, D. A., & Shallice, T. (1986). Attention to action. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation. Advances in research and theory.* New York and London: Plenum Press.