Introduction

In which we give a short overview of this dissertation and discuss the underlying theories and applied methodologies.

Parts of this chapter were adapted from Chapters 2, 4, & 6 of this dissertation, and from:

Borst, J.P., Taatgen, N.A., & Van Rijn, H. (in preparation). Using Cognitive Architectures to Analyze fMRI Data of Complex Tasks. In A. Johnson & R. W. Proctor (Eds.), *Neuroergonomics: Cognitive neuroscience approaches to human factors and ergonomics*: Palgrave Macmillan.

1

Chapter

Chapter 1

Introduction

Introduction

This dissertation was partly written using an application called Concentrate[1]. Concentrate is not a normal program, it does not let you write, draw, or send emails. No, it actually does not *do* anything. What it does is the opposite: it prevents you from doing too many things at the same time, it prevents you from too much multitasking. While it is often said that our "modern world is a multitasking world" (e.g., Salvucci, Taatgen, & Borst, 2009, p. 1), it has become more and more clear that multitasking is not necessarily a good thing. Many studies have shown that while we spend much of our time performing multiple tasks at the same time (e.g., Carrier, Cheever, Rosen, Benitez, & Chang, 2009; González & Mark, 2004), in general this leads to a decrease in performance (e.g., Alm & Nilsson, 1994; Brookhuis, de Vries, & de Waard, 1991; Gillie & Broadbent, 1989; Monk, Trafton, & Boehm-Davis, 2008). In fact, my promotor once said that one of the things that could explain the success of great scientists is their ability to *mono*task, their ability to concentrate on one thing at a time. That is where the aptly named program Concentrate comes into this thesis: it forced me to monotask, and thereby enabled me to finish this dissertation.

The topic of this dissertation is human multitasking, and in particular the so-called *problem state bottleneck*: one of the reasons why multitasking is often counter-productive. While in general humans are extremely good at multitasking – when do we truly do one task at a time? – in certain situations our ability to multitask breaks down. On the one hand, multitasking is obviously limited by physical constraints. We simply cannot look at two things at the same time, as texting cyclists prove daily during my ride to work. More interestingly, there are also limitations in our cognitive system that hinder multitasking. One of these is the problem state bottleneck: a limitation in processing intermediate representations that are necessary for a task. To give an everyday example, imagine you go to the living room to pick up John Anderson's latest book. You store the intermediate representation for this task – pick up Anderson's book – in your *problem state resource*, and walk to the living room. On the way, a friend calls you to ask where you are going to have dinner tonight. After finishing the call you find yourself standing in the living room, without the faintest notion about what you were going to pick up there. According to the theory that I will present in this dissertation, this is caused by a limitation in processing intermediate representations in our brain: the problem state bottleneck.

To support the idea of a problem state bottleneck, I will present several experiments and a computational theory of how intermediate representations are processed in our minds. As support for this theory, I will not only look at behavioral data, but also relate the theory to neuroimaging data. However, before turning to the problem state bottleneck, I will first give a short overview of existing multitasking theories. Multitasking has been investigated for over a century, and especially the threaded cognition theory (Salvucci & Taatgen, 2008, 2011) is of great importance for the current work. This theory will therefore be discussed in some detail below. Furthermore, both threaded cognition and

---

[1] *http://getconcentrating.com/*

the models in this dissertation were implemented in the cognitive architecture ACT-R (e.g., Anderson, 2007). I will therefore also briefly introduce cognitive architectures and ACT-R, followed by a discussion of how cognitive architectures can be combined with neuroimaging research. I will end this introduction with an overview of the other chapters.

Multitasking Theories

As early as 1931, Telford investigated interference due to human multitasking. He introduced the psychological refractory period (PRP) paradigm, and showed that people are slower to respond to the second of two tasks when these tasks have to be performed concurrently. Since Telford, many theories have been put forward to explain interference effects in multitasking (see for overviews, Meyer & Kieras, 1997a; Salvucci & Taatgen, 2008). Theories on multitasking can be divided into three general groups: bottleneck theories, resource theories, and cognitive control theories. Bottleneck theories assume fixed bottlenecks in human cognition that can only process one task at a time, causing interference when used by multiple tasks concurrently (e.g., Broadbent, 1958; Keele, 1973; Pashler, 1994; Welford, 1952). Theorists have identified several different bottlenecks, ranging from perceptual bottlenecks (e.g., Broadbent, 1958), to response-selection bottlenecks (e.g., Pashler, 1984; 1994), to motor bottlenecks (e.g., Keele, 1973). To unify these different bottleneck accounts, resource theories were introduced. These theories assume that attention can be flexibly employed, and that multitasking interference occurs when cognitive resources are required by multiple tasks at the same time, but not when tasks require different resources (e.g., Kahneman, 1973; Navon & Gopher, 1979; Wickens, 1984, 2002). A third research tradition focuses on executive processing and cognitive control to explain multitasking interference (e.g., Baddeley, 1986; Cooper & Shallice, 2000; Meyer & Kieras, 1997a, 1997b; Norman & Shallice, 1986). In these theories, multitasking interference arises because of scheduling problems between tasks. That is, while tasks could in principle be carried out concurrently, executive control mechanisms enforce a certain task order, leading to interference. Using a cognitively bounded rational analysis, Howes, Lewis, and Vera (2009) have recently shown that to best account for at least the classical PRP effect (Schumacher et al., 1999; Telford, 1931) a theory needs cognitive control mechanisms, a motor bottleneck, and a response-selection bottleneck.

Based on the large body of data collected since the 1930s, detailed computational cognitive models of multitasking have been developed, ranging from concurrent multitasking (e.g., Kieras, Meyer, Ballas, & Lauber, 2000; Salvucci, 2005) to task switching (e.g., Altmann & Gray, 2008; Gilbert & Shallice, 2002; Monsell, 2003; Sohn & Anderson, 2001) to sequential multitasking (e.g., Altmann & Trafton, 2007; Salvucci, Monk, & Trafton, 2009). These computational models make it possible to predict the amount of interference between tasks on a quantitative level. Threaded cognition, a recent theory of human multitasking, combines all elements above and was implemented as a computational model (Salvucci & Taatgen, 2008, 2011). In combination with the cognitive architecture ACT-R (Anderson, 2007), it made a specific prediction that is the basis of this dissertation: it predicted the problem state bottleneck.

Threaded Cognition's Prediction

Threaded cognition is a general theory of human multitasking (Salvucci & Taatgen, 2008, 2011; Salvucci, Taatgen, et al., 2009). It assumes multiple different bottlenecks, and states that while multiple tasks can be performed concurrently, every resource in human cognition can only process one task at a time and therefore acts as a bottleneck when required by multiple tasks concurrently (cf. Byrne & Anderson, 2001). Depending on the requirements of the tasks at hand, these bottlenecks lead to different patterns of interference. Thus, the key assumption of threaded cognition is that although *several tasks* can be active at the same time, a particular *resource* can only be used by a *single task* at a time.

For instance, if two tasks want to use the visual system at the same time, only one of them can proceed and the other task will have to wait. In the case of the visual system this is quite obvious: we can only look at one object at a time. However, the same mechanism is assumed to hold for more central resources such as memory. According to threaded cognition, if two tasks want to retrieve a fact from memory at the same time, only one of them can proceed and the other task will have to wait. On the other hand, no interference is predicted if one task uses the visual system while another task retrieves a fact from memory. Thus, as long as the resource requirements of the different tasks do not overlap in time, threaded cognition predicts no interference, but as soon as a particular resource is concurrently needed by two or more tasks, that resource will act as a bottleneck and delay the execution of the combined process. This aligns with the intuition that if two tasks require the same cognitive constructs, the tasks will interfere (e.g., talking and reading both require our language faculties, while talking and walking require different resources).

It was shown that threaded cognition could account for interference caused by two peripheral bottlenecks (vision, motor) and two cognitive bottlenecks (procedural and declarative memory; Salvucci & Taatgen, 2008). In addition, based on its integration in the cognitive architecture ACT-R (e.g., Anderson, 2007), one more source of multitasking interference was predicted: the so-called problem state resource[2]. The problem state resource is used to maintain intermediate representations that are necessary for performing a task. For example, when calculating '2 + 3 × 4' mentally, one might use the intermediate representation '2 + 12'. According to the ACT-R theory, only a single intermediate representation can be maintained at a time, which should lead to interference when multiple representations are required concurrently.

Previously, we have presented results (Borst & Taatgen, 2007) that illustrated the potential role of the problem state resource as a bottleneck in multitasking. In that study, participants had to enter an address in a simulated navigation device while driving a simulated car. Both tasks had two versions: one that required maintaining

[2] *The* imaginal buffer *in ACT-R terminology.*

intermediate representations, and one in which there were no intermediate results. When both tasks required an intermediate representation, performance was slower and more error-prone than could be explained by the difficulty of the separate tasks alone, indicating a bottleneck in processing intermediate representations. However, the setup of that study was relatively under-constrained, making it difficult to derive precise conclusions. In this dissertation I will build on these results, and develop more precise experiments to investigate the problem state bottleneck. To account for the results of these experiments, I will present cognitive computational models that were implemented in the cognitive architecture ACT-R, and validate those models using neuroimaging data. I will now give a brief overview of these methodologies.

Methodologies

Cognitive Modeling & Cognitive Architectures

In 1973, Allen Newell boldly argued that psychology focuses too much on isolated tasks, and as a result does not progress much beyond solving 'small questions': He was worried that psychology would never integrate the results of the many separate experiments into a unified theory of human cognition (Newell, 1973). As a solution, he proposed cognitive architectures: unified theories of cognition in which computational processing models can be developed for a wide variety of tasks. The use of computational models forces one to specify theories at a very precise level, while developing different models within one theory ensures that models do not explain isolated phenomena, but that basic mechanisms are shared between tasks. Currently, there are several cognitive architectures in development, for instance Newell's SOAR (Newell, 1990), EPIC (Meyer & Kieras, 1997a), and ACT-R (Anderson, 2007).

Following Newell's suggestion, in this dissertation we implemented all models in the cognitive architecture ACT-R (Anderson, 2007). This means that the underlying resources of the presented models were validated previously, and could now be used to investigate our theory of the problem state bottleneck (e.g., Cooper, 2007; Newell, 1990). Moreover, because the architecture specifies how humans move the mouse or retrieve a fact from memory, we could develop models of complete tasks (in contrast to single mechanisms), enabling a direct comparison between human and model data. This is especially important for models of multitasking behavior, in which the interaction between cognitive and peripheral resources often causes the observed behavior (Kieras & Meyer, 1997; Van Maanen, Van Rijn, & Borst, 2009). In the next section we will introduce ACT-R, and explain how it maps onto the multiple resource theory of threaded cognition.

ACT-R

Figure 1.1 shows an overview of the ACT-R architecture (Anderson, 2005, 2007; Anderson, Bothell, et al., 2004). ACT-R assumes that the human cognitive system can be described as a system of largely independent modules (cognitive resources)

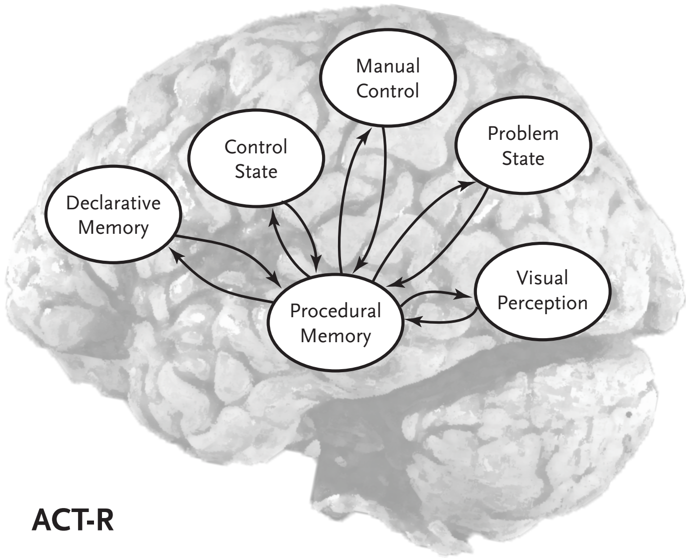

Figure 1.1 Core modules of the ACT-R cognitive architecture.

that interact through a central production system. It can perceive outside information through its visual module and its aural module (not shown in Figure 1.1), and act in the world through its manual module, which operates 'the hands' of ACT-R. To store information it uses a declarative memory store and a procedural memory store, while the control state maintains the current goal of the model. The problem state resource is used to store intermediate representations of a task, and is the main interest of this dissertation. According to our theory, it can only maintain at most one representation at a time, and therefore acts as a bottleneck in multitasking. In general, threaded cognition assumes that every ACT-R module constitutes a bottleneck: all modules can only proceed in a serial fashion, and therefore cause multitasking interference when required by multiple tasks concurrently (Byrne & Anderson, 2001; Salvucci & Taatgen, 2008, 2011).

Note that Figure 1.1 only shows the core modules of the architecture (the aural and vocal module are not shown). Other (or alternative) modules have also been developed, for instance to account for timing (Taatgen, Van Rijn, & Anderson, 2007; Van Rijn & Taatgen, 2008), blending in declarative memory (Lebiere, Gonzalez, & Martin, 2007), and robot perception (Trafton, Bugajska, Fransen, & Ratwani, 2008). Furthermore, while the ACT-R architecture mainly functions on a relatively high level (Newell's cognitive and rational bands, Newell, 1990; see also Anderson, 2002), of many modules more detailed lower-level versions have been proposed, for instance for declarative memory (Van Maanen, Van Rijn, & Taatgen, in press), procedural memory (Stocco, Lebiere, & Anderson, 2010), and visual perception (O'Reilly & Munakata, 2000; Salvucci, 2001).

Model-Based Neuroimaging

For a long time, the field of information processing psychology that argued for cognitive architectures essentially ignored the brain (Anderson, 2007). For instance, Newell states in 1980 that "symbolic behavior (and essentially rational behavior) becomes relatively independent of the underlying technology. Applied to the human organism, this produces a physical basis for the apparent irrelevance of the neural level to intelligent behavior." (Newell, 1980, p. 175). However, since the 1990s cognitive psychologists recognize the importance of the system in which intelligence is realized, and started connecting cognitive architectures to neuroimaging data. Anderson made this very explicit in his definition of a cognitive architecture in 2007 (p. 7): "A *cognitive architecture* is a specification of the structure of the brain at a level of abstraction that explains how it achieves the function of the mind."

One of the reasons for connecting cognitive architectures to neuroimaging data is that many models have a complexity that cannot be fully justified on the basis of behavioral measurements alone (e.g., Myung, 2000; Pitt & Myung, 2002; Roberts & Pashler, 2000). That is, there are so many degrees of freedom in developing a model that models are often under-constrained by behavioral data. To strengthen the constraints on cognitive models that are developed in the cognitive architecture ACT-R, a methodology was developed for mapping model activity on brain activity (for a concise explanation, see Anderson, Fincham, Qin, & Stocco, 2008). This way, models are not only constrained by behavioral data, but also by neuroimaging data. Figure 1.1 roughly shows the mapping between ACT-R's modules and the brain (for details, see Chapter 4 and 5).

Usually, the connection between ACT-R's modules and fMRI data is implemented using predefined Regions-Of-Interest, which provide a mapping between the brain and components of the architecture (Anderson, 2007; Anderson et al., 2008). For instance, activity in the motor resource of ACT-R should correspond to neural activity in a predefined region in the motor cortex (see Figure 1.1). We applied this methodology in Chapter 4 of this dissertation to investigate whether our model made plausible fMRI predictions.

More recently, a new methodology has emerged in fMRI research: model-based fMRI (e.g., Gläscher & O'Doherty, 2010; O'Doherty, Hampton, & Kim, 2007). This analysis technique shows regions in the brain where neural activity significantly correlates with model activity. In Chapter 5, we applied this technique for the first time to a model developed in a cognitive architecture. Model-based fMRI is a promising new method, because it gives a functional explanation of fMRI data by directly linking the data to model constructs (which naturally perform required functions of the model). Functional neuroimaging, especially fMRI, has often been criticized of not contributing anything significant to our understanding of the mind as there is no direct mapping between data and function (e.g., Coltheart, 2004; Coltheart, 2006; Fodor, 1999; Harley, 2004; Page, 2006; but see e.g., Aue, Lavelle, & Cacioppo, 2009; Friston, 2009; Hagoort, 2008; Henson, 2005, 2006; Jonides, Nee, & Berman, 2006; Logothetis, 2008). As Fodor put it forcefully: "If the mind happens in space at all, it happens somewhere north of the neck", on which basis he discounted most fMRI research from explaining anything about our cognitive system (Fodor, 1999). However, especially by combining model-based fMRI with models grounded in a cognitive architecture (which we have shown is possible in Chapter 5), we can at least partly avoid these criticisms, and use fMRI to learn more about the functioning of our cognitive system.

Overview of this Dissertation

As stated above, this dissertation is about the problem state bottleneck. I will present behavioral, model-based, and neuroimaging support for the existence of this bottleneck. First, I will present three behavioral experiments and accompanying cognitive models in Chapter 2. These experiments provide initial support of a problem state bottleneck. In Chapter 3, the behavioral support is extended with an experiment that shows how the problem state bottleneck can be bypassed. In addition, pupil dilation data will be presented in Chapter 3, to show that the bottleneck is associated with an increase in mental workload. In Chapter 4 and 5, I turn to neuroimaging data to validate the cognitive model that was presented in Chapter 1. First, in Chapter 4, a region-of-interest analysis is applied to test *a priori* predictions of the cognitive model. Second, in Chapter 5, the novel model-based fMRI analysis technique is used to show where in the brain the different resources of the model are most likely represented. To conclude, in Chapter 6, our final theory of how intermediate representations are processed in the mind will be presented. This theory will be backed up with data presented in the other chapters of this thesis, and with data of two new behavioral experiments.