Summary & Concluding Remarks

In which we briefly summarize the findings presented in this dissertation and finish with some concluding remarks.

7

Chapter

Chapter 7

Summary & Concluding Remarks

Summary & Concluding Remarks

I started this dissertation discussing an application that makes it harder to multitask: Concentrate[1]. According to the application's website, by enforcing monotasking it "helps you work and study more productively". A possible reason why Concentrate might make you more productive is the topic of this dissertation: the problem state bottleneck. In the preceding chapters we have shown that, due to this bottleneck, having to maintain multiple intermediate representations at the same time leads to a decrease in performance, both in time and accuracy. By focusing on a single task – for example with the help of Concentrate – it is more likely that you use at most one intermediate representation, resulting in better overall performance. In this last chapter, I will briefly summarize our findings and the resulting theory, and along the way discuss some high-level implications (more detailed discussions of the results can be found in the previous chapters). This chapter is organized around the different methodologies that we applied: behavioral experiments, cognitive modeling, pupil dilation, and neuroimaging.

Behavioral Results

First evidence for a problem state bottleneck came from the three behavioral experiments that we discussed in Chapter 2. In the first experiment, by varying the use of intermediate representations and thereby the use of the problem state resource, we showed that performance decreased considerably when more than one representation was required at a time. The second and third experiment were carried out to show that this decrease in performance was not caused by an effect of memory load or by a phonological loop bottleneck, respectively. In Chapter 3, we extended the support for a problem state bottleneck by showing that the interference was indeed due to the use of intermediate representations: when such a representation was presented in the environment, performance suffered less than when a representation had to be maintained mentally. These results run counter to classical working memory theories that assume that people can maintain up to 7±2 items in short-term memory (Miller, 1956). However, they match more recent theories, which propose a very limited focus of attention in working memory of only one or two items that can be used without a time cost (e.g., Cowan, 1995; Garavan, 1998; Jonides et al., 2008; McElree, 2001; Oberauer, 2002, 2009).

Following these results, in Chapter 6 we took a closer look at what happens to intermediate representations that cannot be maintained in the problem state resource. According to the underlying theories, ACT-R and threaded cognition, these representations are stored in a declarative memory store in which their memory strength decays over time. To test this assumption, we conducted two so-called interruption experiments. In these experiments, a primary task was interrupted for varying durations by a secondary task. The data showed that (1) when both tasks needed

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

an intermediate representation the time to resume the primary task was higher than in the other conditions, and (2) that in that condition the time to resume the primary task increased with interruption duration. This matched the predictions of our theory, and indicated that intermediate representations are indeed stored in a declarative memory store when they cannot be maintained in the problem state resource.

Cognitive Models

To account for the results of the experiments we developed computational cognitive models. These models were instantiated in the cognitive architecture ACT-R (e.g., Anderson, 2007). To simulate the multitasking aspects of the experiments we used threaded cognition theory (Salvucci & Taatgen, 2008, 2011). On the one hand, the models show that the observed results can indeed be explained by a bottleneck in the problem state resource. On the other hand, the modeling results also cross-validate the ACT-R theory and threaded cognition. First, our models add new tasks to the expanding set of data ACT-R can account for (see also http://act-r.psy.cmu.edu/), and show thereby that ACT-R is a plausible psychological theory and at the same time a useful modeling framework. With respect to threaded cognition, the modeling accounts show that a multiple-bottleneck theory can explain our datasets, while this would have been difficult with a single central bottleneck account (e.g., Pashler, 1994). For example, while we mostly focused on the problem state bottleneck, to explain the effects of the listening task (Experiment 3, Chapter 2) the bottleneck in declarative memory was crucial. This also shows the added value of using a cognitive architecture: the different effects of, and the interactions between, the problem state resource (Experiment 1 & 2, Chapter 2), the visual resource (Chapter 3), and declarative memory (Experiment 3, Chapter 2; the interruption experiments, Chapter 6) were all necessary to explain the data. The effects of the different resources automatically follow from using a cognitive architecture, because the resources were implemented and supported previously (Anderson, 2007; Cooper, 2007; Kieras & Meyer, 1997; Newell, 1990; Van Maanen et al., 2009). If one would want to model our datasets with 'single-issue models', many more ad-hoc assumptions would have to be made, resulting in weaker modeling accounts.

Pupil Dilation Results

To investigate whether the decrease in performance due to the problem state bottleneck was accompanied by an increase in mental workload, we measured people's pupil dilation in Chapter 3. In addition to the interaction effects in the response time and accuracy data, we indeed also observed an over-additive interaction effect in pupil size, with most dilated pupils in the conditions where more than one intermediate representation had to be used. As it is well known that pupil dilation reflects mental workload (for reviews, see e.g., Beatty, 1982; Steinhauer & Hakerem, 1992), we interpreted this as evidence for increased mental effort when the limits of the problem state resource were reached. While it is still unclear what exactly is reflected by the pupillary response, these results could indicate that one factor is the use or maintenance of intermediate representations.

Neuroimaging Results

As the behavioral, modeling, and pupil dilation results all indicated a problem state bottleneck, the next step was to locate the neural correlates of this bottleneck. To this end we conducted two fMRI experiments (Borst et al., 2009, and Chapter 4 & 5). To analyze the fMRI data we first applied a Regions-of-Interest technique that is commonly used in combination with ACT-R (e.g., Anderson, 2007; Anderson et al., 2008). We used this method to see whether our model was able of making *a priori* predictions of the fMRI data in various regions in the brain, to thereby validate the model (Borst et al., 2009 and Chapter 4). In general the model's predictions were accurate, indicating that the model captures important aspects of the data. However, a discrepancy between the model's predictions and the data was found in the posterior parietal cortex, a region that is associated with the problem state resource. There are at least two possible reasons for this discrepancy: the model's assumptions could be incorrect or the ACT-R–brain link could be incorrect or incomplete.

Because the behavioral data in combination with the model delivered strong support for the problem state bottleneck, we focused on how we could improve the ACT-R–brain link. First, based on the broader fMRI literature and an unpublished dataset (Kao & Anderson, personal communication), we hypothesized that the area in the parietal cortex not only reflects problem state actions, but also visual-spatial activity. When we added this hypothesis to the model, the model-data match improved considerably (Borst et al., 2009). As this is based on a single experimental paradigm it cannot yet be used to validate other models. However, it does warrant deeper investigation, to further improve the mapping between ACT-R and the brain.

The results above are all based on the existing mapping between ACT-R's resources and brain areas. However, this naturally means that the results depend on the correctness of this predefined mapping. To investigate where the resources of our model are best represented in the brain – and to see whether that corresponds to the predefined mapping – we applied a novel model-based fMRI analysis technique in Chapter 5. This analysis technique, in which we regressed the model's predictions directly against the fMRI data, showed that the predictions of the problem state resource mapped best onto a region in the posterior parietal cortex that is slightly anterior to ACT-R's predefined region (in general, the results were surprisingly consistent with the ACT-R mapping). The measured BOLD response in this region was much closer to the model predictions than the response in the predefined region. However, the fit was still not perfect, meaning that either the model has to be improved, or the connection between the model and the brain (for example by adding visual-spatial activity, as implied above). One way to investigate this would be by entering all model resources in a single linear model instead of in several separate linear models (as we did in Chapter 5). By using one linear model, it is possible to have a combination of model resources predict

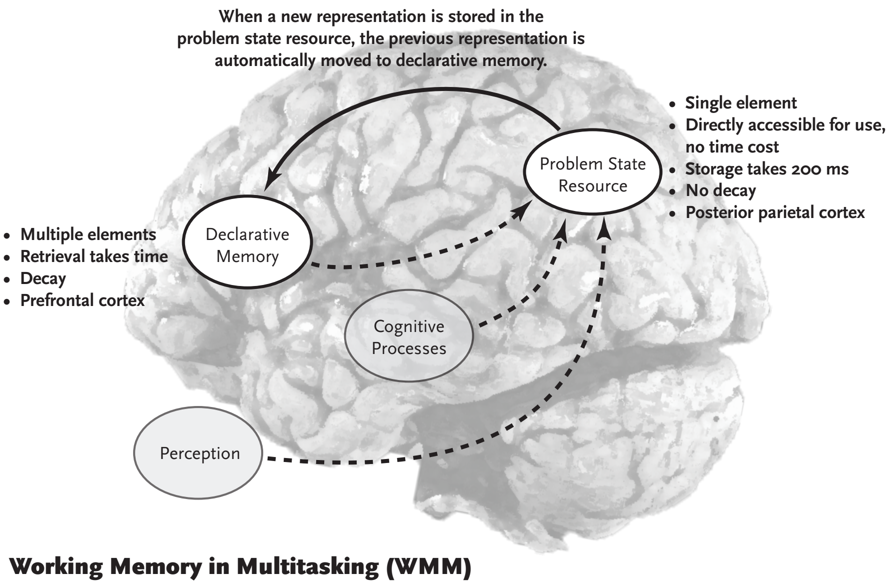

Figure 7.1 Overview of the Working Memory in Multitasking theory.

activity in one brain area. However, for that analysis to work we first need a better experimental paradigm, which dissociates all different resources of the model.

Besides yielding a better mapping between ACT-R and the brain, the model-based fMRI analysis method also leads to more precise function-brain mappings than traditional fMRI analysis methods. Traditionally, fMRI data is regressed against the different conditions of an experiment. The results of those regressions are then subtracted from each other, in principle isolating a certain cognitive function. However, by using a cognitive model – especially one developed in a cognitive architecture – this link between function and data is much more direct, as now a very well described (computationally implemented and previously validated) function is directly regressed against the brain data. While this does not completely solve the problem of fMRI data only showing that "the mind happens [...] north of the neck" (Fodor, 1999), the combination of computational models and fMRI data at least brings us one step closer to understanding the mind.

Conclusion: Working Memory in Multitasking

In Chapter 6 we introduced our over-arching theory: Working Memory in Multitasking (WMM). As this theory is based on all results presented in this dissertation, you could say that it is *the* conclusion of this dissertation. In that sense, Figure 7.1 is *the* summary of this dissertation. As shown in Figure 7.1, the core of the WMM theory is a single-sized problem state resource, which leads to interference when it has to be used by multiple tasks at the same time. When an intermediate representation is removed from the problem state resource it is automatically stored in declarative memory, where it starts to decay. This means that the longer a representation cannot be used in a multitasking situation, the harder it will be to retrieve it and resume an interrupted task. To support the WMM theory we have presented behavioral (Chapter 2 & 3), pupil dilation (Chapter 3), fMRI (Chapter 4 and 5), and computational cognitive modeling support (all chapters). To me, this is the essence of cognitive science: using formal methods in an interdisciplinary manner to investigate and understand the human mind.