# Summary & Conclusion

*In which I summarize the findings and discuss why precision matters.*

6

Chapter

# Chapter 6

# Summary & Conclusion

## Summary & Conclusion

Newell (1973) closed his 20 questions paper with the words: "Maybe all is well, […], and when we arrive in 1992 […] we will have homed in to the essential structure of the mind." (p. 306). Have we, now in 2011, reached this goal? Undeniably, progress has been made since the 70s. The use of computational models has increased the precision of theoretical predictions and cognitive architectures like ACT-R provide a promising tool towards understanding "the essential structure of the mind". Yet, underspecified verbal models and simplifying dichotomies still enjoy great popularity.

Currently, a vast amount of research centers on the debate of whether reasoning and decision making is based on implicit, automatic, and high-capacity, or on explicit, deliberate, and low-capacity processes (e.g., Dijksterhuis, Bos, Nordgren, & Van Baaren, 2006). Another popular dichotomy can be found in the decision sciences, where there is an ongoing discussion whether decisions can better be described by simple non-compensatory heuristics or by more complex compensatory strategies (see e.g., the special issues by Marewski, Pohl, et al., 2010; Marewski, et al., 2011a, 2011b). While tests of binary oppositions can certainly lead to interesting insights, they are, as Newell warned, not always useful. Often, the apparently opposing aspects even represent "two sides of the same coin". The real challenge for understanding cognition lies, therefore, not so much in testing opposing aspects against each other, but in understanding their respective contribution and interaction. To do so, we need to understand the underlying cognitive processes. In this thesis I have shown how the precision provided by computational models can help us to meet this challenge.

## Memory Activation in Diagnostic Reasoning

The starting point of this dissertation was the idea that automatic memory processes can facilitate diagnostic reasoning by providing the reasoner with an adaptive subset of potential hypotheses from memory. The work presented in Chapters 2, 3, and 4 not only supports this idea, but also identifies and tests potential underlying memory mechanisms.

In Chapters 2 and 3 we tested whether and how observed symptoms can activate associated explanations from memory. Using the cognitive architecture ACT-R (Chapter 2) and connectionist constraint satisfaction models based on ECHO (Chapter 3), we implemented several computational models. The models shared the assumption that observed information can activate associated knowledge, but they differed with respect to how the sequentially made observations affected memory activation over time. The results of the models were compared to human data from two behavioral experiments in which we used a probe reaction task to track the availability of different explanations during a sequential diagnostic reasoning task. The basic results were consistent over both approaches: Comparing the probe reaction data to the models' results suggested that the availability of explanations in memory indeed varied as a function of the observed symptoms over time. Furthermore, the probe reaction data was best fit by models in which the influence of observed information did not vary as a function of the number of observations (Chapter 2) and remained stable over time (Chapters 2 and 3).

Both modeling approaches increased the precision compared to mere verbal theories and supported our assumptions about memory activation. However, the approaches differed in the scope of their interpretability. The connectionist models showed how observed data could activate associated knowledge in a network. But how is such a network constructed and which aspects of memory does it reflect? Using ACT-R, and thereby adhering to the constraints set by the underlying memory theory, allowed for interpreting our results more functionally, in terms of general memory mechanisms. We concluded that observed symptoms that are currently in the focus of attention regulate the availability of associated explanations in long-term memory by spreading activation to these explanations. This component of memory activation reflects the usefulness of explanations in the current context.

In Chapter 4, we investigated how the influence of the current context interacts with a second factor that has been proposed to influence an item's availability in memory: its past usefulness (e.g., Anderson, 2007; Thomas, et al., 2008). We conducted a behavioral experiment in which both factors, past and present usefulness, were independently manipulated by a secondary task. Results of this experiment were compared to the predictions of an ACT-R model that was constructed based on our findings in Chapter 2. Participants' performance showed effects of both manipulations, as predicted by the model, suggesting that the past *and* the present usefulness determine the availability of diagnostic hypotheses in memory.

Why is it important to understand memory processes underlying diagnostic reasoning in so much detail? As, for example, discussed by Dougherty et al. (2010), many theories assume that "whatever takes place in the memory system is irrelevant to understanding judgment and decision-making behavior" (p. 337). Our results illustrate why such a simplifying assumption is short-sighted. By using a precise account of memory activation, which was based on general findings about memory mechanisms, we could show how taking into account the contribution of memory processes can lead to a better understanding of hypothesis generation. Such an understanding is not only interesting from a theoretical perspective, but it can, potentially, help to improve real-world decision making.

Consider, for example, the medical setting, where a doctor's ability to generate correct diagnoses can be of vital importance for a patient. The medical literature describes various pitfalls that frequently occur in this setting (for an overview see e.g., Klein, 2005). An understanding of the underlying cognitive mechanisms can help to develop training programs that might reduce such pitfalls. Take, for example, the representativeness heuristic (Tversky & Kahneman, 1974), where diagnosticians strongly rely on information available in the current context (e.g., a patient's symptoms) but seem to ignore information about the base-rate likelihood of the potential diagnoses. Our results suggest that this effect might be due to insufficient personal experience with the diagnoses' base rates. If real-world base rates are not represented in memory in terms of experienced frequencies, the past-experience component of memory activation cannot correctly reflect the real-world base rate information. Consequently, memory activation can be dominated by the current-context component, resulting in behavior as described by the representativeness heuristic. Teaching base-rates in terms of natural frequencies rather than as abstract percentages (see e.g., Sedlmeier & Gigerenzer, 2001, for how that could be done) might help to reduce this pitfall, because it would allow for memory activation to adaptively provide the diagnosis that is most likely not only based on the current context but also based on past experience.

## Decision Making Based on Information from Memory

Whereas in Chapters 2 to 4 we investigated how memory activation affects the availability of information in memory as a function of the past and present environment, in Chapter 5 we investigated how reasoners make decisions by exploiting the availability of memory contents. This investigation is directly related to the second dichotomy mentioned above: the question whether decisions can better be described by simple non-compensatory heuristics or by more complex compensatory decision making strategies. A non-compensatory heuristic has, for example, been proposed in terms of the recognition heuristic (Goldstein & Gigerenzer, 2002). This heuristic states that if only one of two alternatives is recognized, the reasoner will rely on recognition to infer the recognized alternative to have higher values on a given criterion, without using additional knowledge about the alternatives. In contrast to such a non-compensatory heuristic, compensatory decision models assume that people use additional knowledge about the alternatives (Glöckner & Betsch, 2008; Lee & Cummins, 2004). For example, when deciding which of two cities, one recognized and one not, is larger, the reasoner would decide for the recognized city according to the recognition heuristic. According to compensatory models of decision making, the reasoner would take into account additional knowledge about the cities' cues (e.g., does it have an airport?) and might, consequently, conclude that the recognized city is smaller.

As we discussed in Chapter 5, the comparison of these apparently opposing decision strategies has been proven difficult, because the strategies are described at varying levels of detail and are often underspecified relative to the empirical data against which they can be tested. ACT-R allowed us to tackle these issues by implementing several apparently opposing strategies within one modeling framework. This implementation required a high degree of precision and took into account the interplay of the decision strategies themself with, for example, perceptual, memory, and motor processes. Comparing the models to behavioral data from Pachur et al. (2008) showed that models that incorporated a combination of supposedly opposing strategies fit the data best. For example, even participants that always decided for the recognized alternative seemed to occasionally retrieve additional knowledge from memory.

The results of Chapter 5 illustrate again how cognitive architectures like ACT-R can be used to dissolve simplifying dichotomies and increase our understanding of detailed cognitive processes. Furthermore, the results highlight the importance of taking into account the interaction of (deliberate) decision strategies with other aspects of cognition, like the availability of information in memory. For instance, in Pachur et al.'s experiments, all knowledge about the alternatives was taught in a very controlled setup, making it, for example, likely that the different cues about an alternative were equally available in memory. In a more natural setting, the availability of different pieces of information in memory will vary as a function of the environment, which can result in different outcomes for the same decision strategy, and might even cause the use of different decision strategies.

## Conclusion

In this thesis I have shown how the precision provided by computational cognitive models can be used to better understand the cognitive processes underlying complex cognition. Our results illustrate why it is often not useful to construct and test binary oppositions, as it is done in many areas of research. We showed why apparently opposing aspects of cognition, like automatic and deliberate reasoning, or non-compensatory and compensatory decision making, should better be understood as complementary components. For me, the discovery of the respective contribution and interaction of these components represents the real challenge in understanding the human mind, and I hope that the work presented in this thesis represents a step towards solving this challenge. Have we, now in 2011, "homed in to the essential structure of the mind"? While undoubtedly progress has been made, I want to close with a quote from John Anderson: "we are still only a little ways into understanding the answer" (Anderson, 2007, p. 239).