8

Discussion and Conclusion

8.1 Summary and discussion

In this thesis we applied the statistical technique of finite mixture modeling to two domains in the field of cognitive development: the conservation of a continuous quantity (chapter 6), and discrimination (shift) learning (chapters 2 to 5). In both domains, we aimed to distinguish and study categorically distinct behavioral modes, and to investigate transitions between these modes in a statistically sound way. We applied mixture models both to multivariate normally distributed data (conservation; Chapter 6), and to discrete data (discrimination shift learning; Chapters 3 to 5). In chapter 2 we used a bivariate mixture model that combined discrete and normally distributed data. The results from this thesis demonstrate that the finite mixture modeling reveals relevant information, and circumvents methodological problems of traditional approaches. Such approaches have involved group averaging, or a classification based on fixed criteria. These are problematic because they can lead to incorrect inferences (see Gallistel et al., 2004; van der Maas & Straatemeier, 2008).

In the models of this thesis, as in all statistical models, parameter identification, i.e., the ability to infer parameter values from data, is an important issue. In chapter 7, we applied a method for establishing model identification to the class of multinomial processing tree (MPT) models (see Batchelder & Riefer, 1999 for a review). The adopted approach to local identification was discussed previously by Catchpole and Morgan (1997) in capture-recapture (multinomial) modeling, and by Bekker et al. (1994) in structural equation modeling. With respect to MPT models, we demonstrate that the method is tractable, and that it yields useful information concerning identifying constraints and model nesting. In the models applied in this thesis, this method was not applicable, because it requires the availability of sufficient summary statistics. Where it was necessary to establish identification in these models, we resorted to less efficient, but otherwise adequate, simulation methods.

The large part of this thesis concerns the investigation of multiple learning modes in the development of discrimination (shift) learning between early childhood and early adulthood (chapters 2 to 5). We integrated techniques such as finite mixture modeling and existing mathematical learning models for discrimination learning. In addition, we developed a new model to investigate the relation between initial and shift learning. In the following paragraphs, we consider the major findings of the studies in the light of the theory and the present literature on discrimination learning.

In chapters 2, 3, and 5, we examined the existence and the characteristics of multiple modes of learning processes in the initial phase of a discrimination learning task, as discussed in the work of Kendler (1979) and Ashby et al. (1998). According to Kendler (1979), a developmental progression takes place between two qualitatively different modes of learning, i.e., from a slow, incremental mode to a rational, hypothesis-testing mode. We found support for certain aspects of this theory. The analysis of the trial-by-trial data of 4- to 20-year-old participants on a discrimination learning task supported the hypothesis of the existence of two distinct modes of learning rather than a continuous increase in learning efficiency with age (chapters 2, 3; Block et al., 1973; Kendler, 1979; Raijmakers et al., 2001). Specifically, a mixture model that included a fast and a slow learning mode (i.e., as components in the mixture) fitted the data relatively better than models of only one mode of learning. Allowing the learning parameter in the single mode models to vary with age did not change this result. The fast learning mode was consistent with a model of hypothesis-testing given an initial set of four hypotheses, and a small probability of a mistake in the learned state. The probability of the more efficient, hypothesis-testing learning mode increased with age, as predicted by Kendler (1979). Although these findings are in concordance with Kendler's theory, other findings were more difficult to interpret.

In contrast to Kendler's prediction of slow incremental learning, abrupt learning from one trial to the next was observed in the slow learning mode. The accuracy, once learning has taken place, does not deviate significantly from accuracy in the more efficient hypothesis-testing mode. The finding of slow abrupt learning is inconsistent with Kendler's theory, with neural network simulations using the cascade-correlation algorithm Sirois and Shultz (1998), and with predictions based on models of implicit learning, as described by Reber (1993). It seems likely that the children who learned slowly used hypotheses in inefficient ways, e.g., by failing to resample hypotheses due to inefficient feedback processing, by sampling irrelevant hypotheses with a higher probability due to dimension preferences, and/or by expanding the pool of hypotheses with inadequate and complicated combinations of basic hypotheses (informal interviews, see also Phillips & Levine, 1975; Kemler, 1978).

In chapter 4, we investigated the existence of different modes of learning a reversal shift, in which reinforcement contingencies are reversed. In addition, the relation between initial learning mode and shift learning mode was examined. The analysis of the trial-by-trial data of the shift learning phase revealed multiple modes of learning, in particular, a slow and a fast learning mode, and a mode involving an immediate reversal of the previous rule. The learning parameters did not differ between age groups. The probability of perseverating in the previously correct rule did not differ between age groups, either. This absence of age differences in perseveration is somewhat unexpected given the significant decrease of the number of perseverative errors in the Wisconsin Card Sorting Task (WCST; Heaton et al., 1993) that is typically found in developmental studies (e.g., Chelune & Baer, 1986). However, the WCST involves an extradimensional shift instead of a reversal shift, which may explain the diverging results. The mixing proportions of the three modes differed between age groups, and depended on the initial learning mode: Given slow initial learning, the probability of slow shift learning decreased with age, i.e., the probability of positive transfer as a discrete improvement in learning efficiency increased with age. Given fast initial learning, we found slow shift learning (i.e., negative transfer) in almost all 4- to 6-year-olds, whereas this was found only in a small minority of the 7- to 13-year-olds, and was completely absent in the 20-year-olds. Besides these age differences in the use of the learning modes, age differences were also found in the accuracy with which the reversed rule was applied in the learned state. The accuracies in the learned state differed significantly between age groups, and were below criterion accuracy in the 4- to 9-year-olds. This was not the case in the initial learning phase. A possible explanation of the reduced accuracy is interference from the previously learned rule, resulting in occasional responses according to the previous rule. This explanation is consistent with developmental executive function theories (e.g., Zelazo & Frye, 1998, Diamond, 2006).

The results from chapters 3 and 4 provided additional evidence for the effect of preferences for a feature or a dimension on discrimination learning in young children (see also Esposito, 1975). In chapter 3, the mixing proportions of the learning modes differed between the brightness and the shape condition, in the younger children. The probability of using the more efficient learning mode was higher in the brightness condition, in which a rule relating to the brightness dimension had to be learned and shape was irrelevant, than in the shape condition, in which a rule relating to the shape dimension had to be learned and brightness was irrelevant. In chapter 4, the vast majority of 4- to 6-year-olds, who supposedly engaged in hypothesis-testing in the initial learning phase, showed slow shift learning. This may indicate that these children were unable to flexibly test hypotheses (and merely had a coincidental preference for the correct rule in the initial phase), or are unable to do so efficiently because of interference from the initially learned rule.

The possible effect of feature preferences was addressed in the study reported in chapter 5. In addition, we examined the relation between the learning mode and the two executive function components working memory and attentional control. The results replicated the finding of a slow, abrupt learning mode, and a fast, hypothesis-testing learning mode (chapters 2 and 3). As found before, the probability of using the fast, more efficient learning mode increased significantly in the present age range of 4 to 14 years. In addition to these two learning modes, a non-learning mode was found in the youngest age group, which seemed mainly attributable to the manipulation of the relevant dimension. Four- to five-year-olds, who showed a feature preference, and were forced to learn a rule involving unpreferred dimension appear not to be able to flexibly test hypotheses, and more than half of these children showed no sign of learning. In contrast, a small group of the 4- to 5-year-olds, who did not show feature preferences may have tested hypotheses efficiently. Working memory and attentional control measures predicted posterior learning mode probabilities, after controlling for age. This suggested that insufficient working memory capacity and an inability to resist interference of irrelevant and conflicting information hampers the use of the efficient hypothesis-testing mode on the DL task in the tested age range.

Ashby et al. (1998) predicted that children's performance suffers in rule-based tasks, as the discrimination learning task. However, they did not specify whether they expected children to engage the implicit system or their immature verbal learning system. As explained in chapter 3, the finding of abrupt slow learning rather than incremental learning does not necessarily imply that learning in the slow mode is not merely based on simple stimulus-response associative learning (Gallistel et al., 2004; Restle, 1965; Ell & Ashby, 2004; Ashby et al., 1998). However, the present finding renders this hypothesis less likely. Although shift learning results have to be interpreted with caution (e.g., Esposito, 1975), the finding of different shift learning modes in chapter 4, conditional on the initial learning mode seems in conflict with a pure stimulus-response learning associative interpretation of the slow learning mode. This holds in particular for the relatively large probabilities of fast reversal shift learning given slow initial learning (probabilities $> .3$ in all groups). Furthermore, the multimodality of shift learning conditional on age group and initial learning mode suggested the possibility of heterogeneity in the initial learning modes. We identified one source of heterogeneity as preferences for a dimension of the stimuli. As mentioned above, the all-or-none model, which provided the best description to the slow learning process, may accommodate several strategies. A multitude of these strategies could be present simultaneously in the group of slow learning children.

Thus, on the basis of the results of chapters 2 to 5, it seems more likely that the slow and non-learning children engage in inefficient hypothesis-testing strategies, which are related to insufficient working memory and limitations in attentional control. As mentioned in chapter 5, van Duijvenvoorde et al. (2008) showed that the brain areas, which are more active following positive than negative feedback in 8- to 9-year-old children, are the same brain areas, which are more active following negative than positive feedback in adults. This finding provides some support for the interpretation of the slow mode as inefficient hypothesis testing, maybe involving a stronger focus on the confirmation of hypotheses than on the falsification. Note, however, that this study involved a different task, which limits the comparison. Furthermore, as described in chapter 5, in the same age group of 8- to 9-year-olds approximately half of the children seemed to use a hypothesis-testing strategy in performing the discrimination learning task. A possible explanation of the empirical results based on reduced effectiveness of negative feedback in combination with dimensional bias was implemented in a neural network simulation by Berkeljon and Raijmakers (2007).

Suppose that both learning modes relied on the same learning system, i.e., the verbal system (Ashby et al., 1998). One would then expect a gradual increase in efficiency, as the neurological structures underlying the verbal learning system mature. This raises the question how distinct modes of learning, as observed in this thesis can arise from gradual maturation. A formal model that is describes qualitative changes in a dependent variable (e.g., mode of learning) related to gradual change in independent variables (e.g., brain maturation) is the *cusp model*, which was discussed in the developmental psychological context by van der Maas and Molenaar (1992).

8.2 Directions for future research

As mentioned in the introductory chapter, Ashby and Ell (2002) proposed a hierarchy of criteria to distinguish between single and multiple systems of learning: These include criteria in the mathematical, psychological, and neurobiological domains. The mathematical models that provided the best description of the different fast mode and the slow mode of learning were not structurally equivalent, as they are two non-nested submodels of a more general model. However, the differences due to the difference in structure are relatively small compared to the differences due to the quantitative differences in parameter values. In our opinion, this result does not fully satisfy the criterion in the first domain, despite the presence of distinct modes. As the present results do not lend themselves to a definite conclusion about the underlying nature of the slow learning process, more research is needed to elucidate this issue. Latent learning modeling of the performance on learning tasks, which involve blank trials, may help to distinguish between these strategies (Phillips & Levine, 1975). However, especially young children's performance is expected to suffer in these tasks of necessarily increased difficulty. Another promising possibility is the extension of the learning models to include information about the stimulus order (Batchelder, 1971).

Furthermore, as we focused exclusively on behavioral data, we cannot address the question whether the behaviorally distinct learning modes of slow learning and hypothesis-testing correspond to neurologically separated learning systems, as described by Ashby et al. (1998). To address this question, we would require a combination of neuropsychological measurements and mathematical modeling. In addition, the existence of categorically distinct behavioral modes within each age group in the age range of 4-14 years underlines the importance of testing the assumption of homogeneity not only in behavioral studies, but also in neuropsychological studies (see Noppeney, Penny, Price, Flandin, & Friston, 2006).

In this thesis, we investigated category learning processes exclusively in one instance of rule-based category-learning tasks, i.e., the discrimination shift learning task. Due to the design of the task it was not possible to distinguish a broader range of different hypothesis-testing strategies in the older children and adults. However, a task involving more dimensions, which would allow to distinguish more hypothesis-testing strategies, would probably have been too difficult for the youngest children in this thesis. That this task is still of relevance to the literature is evident given that it features as a benchmark for neural network simulation studies of (the development of) human categorization learning and shift learning behavior (e.g., Sirois & Shultz, 1998; Kruschke, 1996; Raijmakers, van Koten, & Molenaar, 1996; Raijmakers, Coffey, Stevenson, Winkel, & Berkeljon, 2009). A good description of the learning processes and their development in humans performing a discrimination shift task is a necessary prerequisite for a reasonable comparison of human and network behavior. In this context, further research is needed concerning learning paradigms other than reversal shift (e.g., extradimensional shifts).

In this thesis, we did not investigate whether the finding of distinct modes of learning generalizes to other category-learning tasks. However, we speculate that different strategies of responding might also underlie the performance on the WCST (Heaton et al., 1993), as used in neuropsychological assessment. As explained in chapter 3, the use of norms in the form of mean scores and standard deviations per age group is problematic in the presence of heterogeneity within age groups. In addition, the finding of below criterion response accuracy in the two youngest age groups (chapter 4) in the learned state of the reversal shift phase highlights the importance of distinguishing between different types of errors in the WCST (as suggested by Barceló & Knight, 2002). Failure to reach the criterion of 10 correct sorts in order to proceed with the next learning phase can be due to the participant's failure to identify the correct sorting rule, or due to a failure to perfectly execute the correctly identified sorting rule. It seems relevant to distinguish between these causes, as witnessed by the large number of complex variables calculated from WCST performance (Heaton et al., 1993). Mathematical modeling may provide a way of distinguishing qualitative and quantitative differences in the performance of healthy and clinical participants, and to measure processes at the individual level (see Bishara et al., 2009).

The elaboration and application of the models has helped to pave the way for an application of related models to several other learning paradigms that involve, for instance, a richer category structure (e.g., Raijmakers, Ryan, Wills, Lea, & Visser, in preparation; Raijmakers, Schmittmann, & Visser, in revision; Haring, Visser, & Raijmakers, 2009). We modified the learning models to permit responses in the form of proportions that follow a beta distribution (rather than discrete responses). This was done in order to study discrimination learning processes in infants, whose proportion of trial-by-trial fixation to the correct stimulus were used as an indicator. As yet, elaboration of these models and of the paradigm is needed to allow successful applications to infant learning experiments. The use of a continuous rather than discrete measurement scale (as in the models for fixation proportions and the models in chapter 6 may allow a more finegrained analysis of the learning processes, and, in particular, a more sensitive test of incremental learning. This continuous measure may be obtained by requiring the participants to indicate a tendency of choice on a continuous scale rather than a forced-alternative discrete choice. Another way is by the inclusion of latency data, e.g., response times, as an additional indicator for the latent states, as we did in a subset of the data in chapter 2. We are currently exploring feedback inspection times (Schmittmann & Raijmakers, in preparation b). Finally, we are exploring the use of hierarchical Bayesian modeling to accommodate continuous individual differences (for instance, in the learning parameters as suggested by Batchelder, 1975).

Concerning the generalization to domains other than learning and conservation, the application of latent Markov models allowed us to distinguish distinct modes of of conditional reasoning, and to gauge the effects of item manipulation on the reasoning modes (Schmittmann & Raijmakers, in preparation a). Apart from the applications in this thesis, heterogeneity in the form of categorically distinct behavioral modes, and transitions between these modes have been detected in a broad range of domains in cognitive development. These include transitive reasoning (e.g., Bouwmeester, Sijtsma, & Vermunt, 2004), conditional reasoning (e.g., Rijmen & De Boeck, 2003), proportional reasoning (e.g., Jansen & van der Maas, 1997), triad classification (Raijmakers et al., 2004), children's knowledge of the physical characteristics of the earth (Straatemeier et al., 2008), and children's understanding that still water is horizontal (e.g., Thomas & Hettmansperger, 2001). Given this widespread presence of categorical heterogeneity and discrete change, we expect to encounter this in many other domains. One task for the future is to elaborate statistical models that can efficiently accommodate such phenomena. This requires an approach that integrates progression in the substantive theory and the development of statistical techniques. One inconvenience in the application of the statistical models used in this thesis is that large sample sizes are required. This inconvenience, however, is a consequence of the latent heterogeneity that is present in the subject of study, rather than of the technique that accommodates the heterogeneity.

8.3 Conclusion

What can we conclude on the basis of the results reported in this thesis? By employing the method of latent mixture Markov modeling, we were able to distinguish categorically distinct behavioral modes in two areas of cognitive development (rule-based category-learning, and conservation of a continuous quantity), and to investigate the transitions between these modes. In the applications to discrimination learning, the results supported the existence of categorically distinct modes of learning, one of which can be interpreted as efficient hypothesis-testing. The probability of using this efficient mode increases with age. In the other mode, we found abrupt slow learning instead of the expected incremental learning. Our interpretation of this slow learning mode is tentative, and we raise doubts concerning its homogeneity. We consider it likely that this mode is due to an inefficient use of the immature verbal learning system (Ashby et al., 1998), which is probably hampered by developmental limitations of working memory and attentional control. The existence of categorical individual differences in cognitive behavior within age groups, as we found in two very different cognitive domains, posits a fundamental problem to the comparison of age groups based on their averages. Mathematical modeling of the behavioral data provides a more demanding, but much better basis for studying development in the presence of such categorical individual differences.