Downloaded from UvA-DARE, the institutional repository of the University of Amsterdam (UvA) http://hdl.handle.net/11245/2.50752

File ID uvapub:50752 Filename 8 Discussion Version unknown

SOURCE (OR PART OF THE FOLLOWING SOURCE):

Type PhD thesis

Title Rules and associations: hidden Markov models and neural networks in the

psychology of learning

Author(s) I. Visser

Faculty FMG: Psychology Research Institute

Year 2002

FULL BIBLIOGRAPHIC DETAILS:

http://hdl.handle.net/11245/1.253151

Copyright

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content licence (like Creative Commons).

In this thesis we investigated rule-following behavior as manifest in implicit learning, and different models of such behavior. We addressed the issue whether neural networks can represent formal language,s like symbolic models in chapter 2 and 3. It emerges from the discussion in chapter 2 that recurrent neural networks and the classical model of cognition, i.e. the syntactic model, share the characteristics of systematicity and productivity. That is, at the level of finite state automata, the classical model and the simple recurrent network perform equally well. For higher order languages, the equivalence is not as well established, although some results suggest that context-free and possibly context-sensitive languages can also be represented in recurrent neural networks (Rodriquez et al., 1999). In fact, the representations that arise in the simple recurrent network while learning regular languages resemble symbolic representations. The main difference between symbolic and subsymbolic representations is that the latter are distributed. This brings with it a host of advantages such as frequency effects in learning, graceful degradation and resistance to noise (Bullinaria, 1999).

In chapters 4 and 5, I presented solutions to statistical problems that arise in fitting hidden Markov models to psychological data. Confidence intervals of estimated parameters are important in assessing whether a fitted model is adequate. In chapter 4, three methods for computing confidence intervals were compared. Likelihood profiling and bootstrapping produced similar results, whereas the finite differences approximation to the Hessian turned out to be inaccurate. Likelihood profiling provides very detailed information about the likelihood function in a neighborhood of the maximum likelihood estimate of a parameter. It can result in skew confidence intervals, which may be of interest in some applications. When compared with likelihood profiling, bootstrapping is easier to implement and computationally less demanding, certainly in the case of large models with many parameters.

Markov and hidden Markov models form a large class of models that have proven to be particularly useful in analyzing learning behavior. Appendix B details the program that I developed for fitting (hidden) Markov models on sequences of categorical data. In chapter 5, statistical issues in fitting hidden Markov models were addressed that are necessary for applications to psychological data. First, I presented a general method of fitting models with equality constraints. Second, model selection criteria were compared using a simulation study. In practical applications, the adjusted BIC criterion is to be preferred. Third, I introduced a goodness-of-fit measure, the prediction error, which is novel in the context of hidden Markov models. In two applications these techniques were illustrated. In particular, a novel way of analyzing generation data from implicit learning experiments is introduced. This analysis allows quantification of subjects' performance in implicit

102 Discussion

learning, which in turn can be to test hypotheses about the nature of sequence knowledge.

8.1 Implicit learning

Implicit learning has been the subject a great deal of research in the past ten years. The most important issue in this field seems to be the status of implicit learning, i.e. does implicit learning really occur? In memory research, which has a much longer history (see Raaijmakers and Shiffrin, In press, for an historical review of memory modeling), there is some consensus as to what is implicit and what is explicit memory. Implicit memory is defined in terms of priming effects, whereas explicit memory is associated with recall and recognition. Such consensus seems to be lacking in the field of implicit learning.

At the outset of implicit learning research, with the seminal paper by Reber (1967), this kind of learning was referred to as unconscious learning. As such, it was contrasted with learning processes in which subjects make a conscious effort to acquire knowledge. Because of this, implicit learning has also been called unintentional or accidental learning. Reber (1967) used verbal reports of subjects to check whether they had any conscious knowledge of the material they had studied. Upon this criterion, this turned out not to be the case.

In recent research, verbal reporting has been replaced by the generation task and the recognition task (e.g. Nissen and Bullemer, 1987; Perruchet and Amorim, 1992; Shanks and Johnstone, 1999). These latter tasks are more comparable with their counterparts in memory research: priming, recall and recognition. Another distinction that is frequently made in the implicit learning literature is that between direct and indirect measures of sequence knowledge. Reaction times are seen as an indirect measure and generation and recognition are seen as direct measures (Jimenez et al., 1996). One could ask however, as a measure of what?

Some researchers claim that reaction times are a measure of implicit knowledge and the knowledge that results from implicit learning is unconscious knowledge (Cleeremans and McClelland, 1991; Reber, 1993). Direct measurements of sequence knowledge are generally interpreted as indicative of explicit knowledge. Hence, other researchers argue, since they found large associations between direct and indirect measurements of sequence knowledge, that sequence knowledge must be (largely) explicit (Perruchet and Amorim, 1992; Shanks and Johnstone, 1999). Logically, however, this is not the only possible interpretation of these results. There are at least two other possible interpretations. First, it can be argued that in sequence learning both implicit and explicit knowledge are acquired. Second, it can be argued that generation performance should be interpreted as an expression of implicit knowledge.

The first interpretation, that both implicit and explicit knowledge are acquired during sequence learning, is at odds with the parsimony principle. In experimental settings, large associations have been found between direct and indirect measurements of sequence knowledge. Hence, suggesting that there are two separate knowledge bases at work in producing these results seems overkeen. On the other hand,

8.1 Implicit learning 103

some dissociations have also been found in experiments (Shanks and Perruchet, In press; Destrebecqz and Cleeremans, 2001), suggesting that there implicit and explicit knowledge have different bases.

The second interpretation, to the effect that generation performance can be interpreted as an expression of implicit knowledge rather than an expression of explicit knowledge, runs counter to our intuitions. That is, in a generation task subjects have to consciously choose which stimulus they think will appear next. If anything may be viewed as explicit, certainly conscious choice making may. However, subjects report that they are not well aware of making such decisions. Rather, they 'go with the flow' and rely on 'intuition' in responding to generation trials. In fact, subjects may be tapping their implicit knowledge in responding to generation trials.

This interpretation can also accommodate the results from post-experimental interviews. Subjects can express only little knowledge in post-experimental interviews. When taken at face value, verbally expressed knowledge simply is explicit knowledge. Some researchers have argued that verbal reporting is not very sensitive in bringing out explicit knowledge. To me this is a strange argument. It leads to a multiplication of explanatory processes. There is, in this view, implicit knowledge and explicit knowledge. On top of that, there is conscious knowledge, which is a part of the explicit knowledge, i.e. the part of explicit knowledge that is consciously accessible.

The results that I presented in chapters 6 and 7 are consistent with the second interpretation above. In both experiments strong associations are found between a direct and an indirect measure of sequence knowledge. In fact, associations were close to one in both cases. Also, for a number of theoretical reasons, I maintain that this interpretation is tenable. First, generation or prediction performance need not necessarily be interpreted as an indicator of explicit knowledge. It is certainly consistent to maintain that anticipation of an upcoming stimulus is guided by implicit processes although the response itself is a conscious act. It is not necessary to suppose that subjects have any reasons for responding in a particular way, i.e. reasons in the Wittgensteinian sense that I discussed in the introduction. Subjects just respond without being able to provide a reason or motivation for doing so. Second, I think it is warranted to take verbal reports at face value. Verbal reports are an expression of consciously accessible knowledge, and that is that. Certainly different task sets may bring out different aspects of knowledge, just as is the case in implicit and explicit memory. However, both forms of knowledge, if indeed there are two forms, seem to be largely unconscious, i.e. not consciously accessible.

In chapter 7, we used quantitative methods to compare verbally reported sequence knowledge and sequence knowledge expressed through the generation task. From the analyses, it emerges that knowledge expressed in the generation task is much more akin to knowledge expressed in the reaction times than to verbally expressed knowledge.

104 Discussion

8.2 Rules revisited

I started this thesis with a discussion of the notion of rule-following behavior. I described the formal notion of rule-following that is part and parcel of the cognitivist tradition (Chomsky, 1980; Fodor and Pylyshyn, 1988). Implicit learning can be interpreted as an experimental paradigm in which the acquisition of rule-following behavior is elucidated. From our analyses using hidden Markov models, it may be concluded that the knowledge that subjects acquire in implicit learning is a form of rule-based knowledge. That is, subjects' behavior is certainly consistent with such an interpretation. On the other hand, the neural network model introduced by Cleeremans and McClelland (1991), seems to suggest that implicit learning is a form of associative learning. In chapters 2 and 3, I argued that the representations in the simple recurrent network are very closely related to the canonical representation of regular languages in finite state automata. Combining these results, the conclusion may be drawn that associatively built representations, as in the simple recurrent network, may be rule-based representations.