

Cognition, 49, (1993) 391-396

Discussion

Connections and symbols: closing the gap

Brian MacWhinney

Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Ling and Marinov (L & M) have constructed an interesting symbolic alternative to current connectionist models of language acquisition. The kind of detailed comparison of connectionist and symbolic models that they are pursuing works to clarify and solidify the basis of modelling as a research tool in cognitive science. Work such as that of Shavlik, Mooney, and Towell (1991) shows that symbolic pattern associators (SPA) can rival artificial neural networks (ANN) across a wide range of tasks. Moreover, these SPA models often run more quickly and can be interpreted through a concrete set of easily summarized production rules. Together, these various properties make SPAs good competitors with ANNs, particularly when issues of psychological and neuropsychological verisimilitude are not at stake.

L & M believe that their model outperforms the connectionist model presented by MacWhinney and Leinbach (1991) (M & L) "by a wide margin". This conclusion is incorrect, because many of the numbers they use to support the conclusion are incorrect. In fact the performance of the two models is remarkably similar along virtually all relevant dimensions. The actual figures that should have been included in Table 4 from L & M are as shown here in Table 1.

The simulation from which the results in Table 1 are taken is a minimalized version of the simulation reported on in M & L. It only uses right-justified representations, one pool of hidden units, and no copy units. These changes were made not to improve performance, but to make it easier for others to replicate these results in their own laboratories. The complete set of files - input files, pattern files, run files, instruction files, and crucial output results in the files bpf.sum.testing and bpf.sum.training - has been available for anonymous FTP in a single tar file since early 1993 from the host poppy.psy.cmu.edu in the /pdp subdirectory. The files can be run using the PlaNet simulator, which is also retrievable over the InterNet through anonymous FTP from boulder.colorado.edu

Table 1. *Training and testing results for Rumelhart and McClelland (R & M), MacWhinney and Leinbach (M & L), and Ling and Marinov (SPA)*

	<i>R & M</i>	<i>M & L</i>	<i>SPA</i>
Verbs in training	420	1200	1038
Verbs types in testing	86	15irreg., 87reg.	360
Per cent correct on			
Total	97.0	99.3	99.2
Regulars	98.0	100.0	99.6
Irregulars	95.0	80.0-90.0	96.6
Per cent correct on testing			
Total	56.9-73.3	80.0	89.4
Regulars	66.7-83.3	91.0	89.4
Irregulars	7.1-21.4	27.0	38.1

¹ At epoch 4200 there were 80% correct and at epoch 24000 there were 90.0% correct. ¹ These are the results for testing at epoch 4200.

in the directory pub/generic-sources/PlaNet. L & M were made aware of these files in early 1993, but apparently ran into technical problems because of their decision to substitute the Xerion simulator for the PlaNet simulator. We recommend using the PlaNet simulator to avoid problems of this type.

The numbers given in the revised version of Table 4 show that the three simulations do equally well during training, but that the M & L and SPA simulations outperform the Rumelhart and McClelland (1986) (R & M) simulation in testing. These numbers refute L & M's basic claim that "the SPA outperformed the connectionist models by a wide margin". In fact, the M & L model narrowly outperforms the SPA model on the training and testing of regulars, and the SPA model narrowly outperforms the M & L model on the training and testing of irregulars. No particular significance should be attached to these narrow differences, since they can be affected by the changes of parameters in either class of models. The results given in the revised version of Table 4 also call into question the claim by Ling and Marinov that the M & L model represents a step back from the R & M connectionist model in overall error rate. This is quite clearly not the case.

L & M also attempt to argue for the superiority of the SPA model on the basis of results from the teaching of novel pseudowords. Here, again, there is good evidence that no important differences exist between the M & L and SPA models. For the simulation on poppy.psy.cmu.edu, there is data in the /outfiles subdirectory for analyses parallel to the Prasada and Pinker (1993) analyses conducted by L & M. The results for the M & L model when tested at epoch 4200 are remarkably close to those for the SPA model. For Figure 1(a) (Figure 3 (a) in L & M), the data points are: prototypical 5, intermediate 2, and distant 1. Thus the M & L model is better than the R & M model, but not as good as the SPA model for Figure 1(a). For the more important results conveyed in L & M's Figure 3(b)




Figure 1. Comparison of human subjects with the Ling and Marinov (SPA), Rumelhart and McClelland (R & M), and MacWhinney and Leinbach (M & L) overgeneralization results for the stimuli studied by Prasada and Pinker (1993). (a) Pseudo-irregular verbs. (b) Pseudo-regular verbs.

(Figure 1(b) here), the M & L model matches the fit provided by the SPA model. The numbers are: prototypical 10, intermediate 6, and distant 8. As in the results for training and testing of real words, the M & L and SPA models are virtually indistinguishable in predictive power and both are more accurate than the R & M model.

L & M complain on page 247 about the lack of "openness" in the M & L reporting of errors for the simulation. We would like to believe that they are not suggesting that our simulation might have produced forms like "membled" which we somehow swept under the rug. This was certainly not the case, if our simulation had produced any bizarre errors, we would have reported them. In fact, our results closely match those of L & M's Table 5 for SPA. The only errors we found were regularizations, no-change errors, and vowel change errors. This is the same type of distribution of errors that L & M report for their SPA model. The details can be found in the file `bpf.e4200.ir_test.miss` on `psyling.psy.cmu.edu`. Here is the complete list: befailed, bided, bleded, forbided, forgived, misleded, overtaked, quitted, steng (from sting), strang (from string), and undloed.

L & M claim that M & L failed to report the results for training on the present participle and the third person singular. In fact they state that "there is no indication that they even produced any results for the past participle, present participle, and third person singular, so their claim that the model has succeeded in learning the English verb paradigm is unfounded". However, the results for these other three cells in the paradigm are clearly stated on page 145 of M & L. The model learned present participle and the third person singular early and without error. By contrast, the results for the past participle were much like those for the past tense, with errors on irregulars persisting until the end of training. See pages 145-147 in M & L for further details.

In a final attempt to demonstrate the superiority of the SPA model, L & M construct some "ultimate comparisons" in their postscript in section 6. Unfortunately, the versions of the M & L model that they create in this section mistakenly used left-justification rather than right-justification. The simulation on `poppy.psy.cmu.edu` crucially relies on its right-justified representation. L & M's use of a left-justified representation deprives their version of the M & L model of the crucial cues it needs for the formation of the English past tense. As Slobin (1973) has noted, it is important for children (and models) to "pay attention to the ends of words". Without right-justified information, the model cannot function properly and this vitiates the various comparisons in the postscript section.

In all relevant regards, the M & L and SPA models are clearly close competitors. The M & L model outperforms the SPA model for regulars, but the SPA model outperforms the M & L models for irregulars. Both closely match human performance on novel pseudowords. Both are able to acquire the full English verb paradigm with few errors on the present participle and the third person singular. There is a very good reason for the equivalent performance of these two models. Both models are attempting to acquire the same target forms. This is true in large part because we provided L & M with their training set. In addition, both used the same set of words from Prasada and Pinker. When two computationally powerful systems are given the same set of input data, they both extract every bit of data regularity from that input. Without any further processing, there is only so much blood that can be squeezed out of a turnip, and each of our systems extracted what

they could. There is no special "magic" in either connectionist nets or symbolic pattern classifier systems that can go beyond the data given.

If one wants to go beyond the basic input data, one has to begin to postulate additional mechanisms. L & M take two excursions down this path. First, they devise a system for modifying the "m" parameter in order to simulate u-shaped learning curves. Although the notion of a parameter that varies across the child's development is nothing new to developmental psychologists, it is not something that can be simply accepted without strong independent supporting evidence. L & M provide no such evidence. Their second excursion down the pathway of additional *ad hoc* mechanisms involves the constructions of compression strategies to reduce low-level rules to true linguistic generalizations. Even if one were to ignore the various epistemological problems these compressions entail, it is not clear that many of the defenders of the symbolic approach would accept this version of a language acquisition device.

Leaving aside the technical problems raised by L & M's attempt to demonstrate clear superiority for their SPA model, one needs to focus on the important positive aspects of their work. Perhaps the most important contribution is the way in which the work of L & M tends to alter the texture of the debate between symbolic and connectionist approaches. As symbolic approaches such as IDS, SPA, analogical modeling (Skousen, 1992), ACT-R (Anderson, 1993), or SOAR (Newell, 1990) move toward massive parallelism and low-level generalizations that are then post-compiled into higher-level generalizations, they look more and more like connectionist networks. At the same time, as connectionist networks incorporate limited modularity (Jacobs, Jordan, & Barto, 1991), standard representational systems (MacWhinney & Leinbach, 1991; MacWhinney, Leinbach, Taraban, & McDonald, 1989), and external control systems (Gupta & Schneider, 1991; Jain, 1992) they start to look more and more like symbolic systems. This type of convergence between supposedly exclusive and antithetical systems suggests that empirical data on language learning for topics such as morphophonology may now be sufficiently powerful to constrain modeling effort into a tight range of conceptual channels. The fact that supposedly very different models end up taking on similar shapes bodes well for the development of cognitive science.

References

- Anderson, J. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Gupta, P., & Schneider, W. (1991). Attention, automaticity and priority learning. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Jacobs, R., Jordan, M., & Barto, A. (1991). Task decomposition through competition in a modular connectionist architecture: the what and where vision tasks. *Cognitive Science*, 15, 219-250.
- Jain, A. (1992). Generalization performance in PARSEC: a structured connectionist parsing architecture. In J. Moody, S. Hanson, & R. Lippman (Eds.), *Advances in neural information processing systems 4*. San Mateo: Morgan Kaufmann.
- Ling, C.X., & Marinov, M. (1993). Answering the connectionist challenge: a symbolic model of learning the past tenses of English verbs. *Cognition*, 49, 235-290.
- MacWhinney, B., & Leinbach, J. (1991). Implementations are not conceptualizations: revising the verb learning model. *Cognition*, 29, 121-157.
- MacWhinney, B., Leinbach, J., Taraban, R., & McDonald, J. (1989). Language learning: cues or rules? *Journal of Memory and Language* 28, 355-377
- Newell, A. (1990). *A unified theory of cognition*. Cambridge, MA: Harvard University Press.
- Prasada, S., & Pinker, S. (1993). Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes*, 8, 1-56.
- Rumelhart, D., & McClelland, J. (1986). On learning the past tenses of English verbs. In Rumelhart, D., McClelland, J., & the PDF Research Group (Eds.), *Parallel distributed processing* (Vol. 2, pp. 216-271). Cambridge, MA: MIT Press
- Shavlik, J., Mooney, R., & Towell, G. (1991). Symbolic and neural learning algorithms: an experimental comparison. *Machine Learning*, 6, 111-143.
- Skousen, R. (1993). Analogy: a non-rule alternative to neural networks. In R. Corrigan, S. Lima, & M. Noonan (Eds.), *The reality of linguistic rules*. San Diego: Academic Press.
- Slobin, D.I. (1973). Cognitive prerequisites for the development of grammar. In C.A. Ferguson & D. I. Slobin (Eds.) *Studies of child language development*. New York: Holt Rinehart, and Winston.