reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–114.
Cowan, N., Chen, Z., & Rouder, J. N. (2004). Constant capacity in an immediate serial-recall task: A logical sequel to Miller (1956). *Psychological Science, 15,* 634–640.
Davis, C. J. (1999). *The self-organising lexical acquisition and recognition (SOLAR) model of visual word recognition* (Unpublished doctoral dissertation). University of New South Wales, Sydney, Australia.
Davis, C. J., & Bowers, J. S. (2004). What do letter migration errors reveal about letter position coding in visual word recognition? *Journal of Experimental Psychology: Human Perception and Performance, 30,* 923–941.
Davis, C. J., & Bowers, J. S. (2006). Contrasting five different theories of letter position coding: Evidence from orthographic similarity effects. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 535–557.
Elman, J. (1998). Generalization, simple recurrent networks, and the emergence of structure. In M. A. Gernsbacher & S. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (p. 6). Mahwah, NJ: Erlbaum.
Farrell, S., & Lewandowsky, S. (2002). An endogenous model of ordering in serial recall. *Psychonomic Bulletin & Review, 9,* 59–79.
Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition, 28,* 3–71.
Grossberg, S. (1978). Behavioral contrast in short-term memory: Serial binary memory models or parallel continuous memory models? *Journal of Mathematical Psychology, 17,* 199–219.
Grossberg, S. (1980). How does the brain build a cognitive code? *Psychological Review, 87,* 1–51.
Grossberg, S., & Pearson, L. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: Toward a unified theory of how the cerebral cortex works. *Psychological Review, 115,* 677–732.
Henson, R. N. A. (1996). *Short-term memory for serial order* (Unpublished doctoral dissertation). University of Cambridge, Cambridge, England.
Hulme, C., Suprenant, A. M., Bireta, T. J., Stuart, G., & Neath, I. (2004). Abolishing the word-length effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30, 1,* 98–106.
Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–517.
Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427–466.
Jefferies, E., Frankish, C. R., & Lambon Ralph, M. A. (2006). Lexical and semantic binding in verbal short-term memory. *Journal of Memory and Language, 54,* 81–98.
Lee, C. L., & Estes, W. K. (1981). Item and order information in short-term memory: Evidence for multilevel perturbation processes. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 149–169.
Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review, 96,* 25–57.
Marcus, G. F. (1998). Rethinking eliminative connectionism. *Cognitive Psychology, 37,* 243–282.
Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999, January 1). Rule learning in seven-month-old infants. *Science, 283,* 77–80.
McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning-systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102,* 419–457.
McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review, 88,* 375–407.
McClelland, J. L., Rumelhart, D. E., & The PDP Research Group. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81–97.
Miller, G. A., & Selfridge, J. A. (1951). Verbal context and the recall of meaningful material. *American Journal of Psychology, 63,* 176–185.
Munakata, Y., & O'Reilly, R. C. (2003). Developmental and computational neuroscience approaches to cognition: The case of generalization. *Cognitive Studies, 10,* 76–92.
Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review, 105,* 761–781.
Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language-acquisition. *Cognition, 28,* 73–193.
Pinker, S., & Ullman, M. (2002). The past and future of the past tense. *Trends in Cognitive Science, 6,* 456–463.
Prasada, S., & Pinker, S. (1993). Generalisation of regular and irregular morphological patterns. *Language and Cognitive Processes, 8,* 1–56.
Thomas, M. S. C., & McClelland, J. L. (2008). Connectionist models of cognition. In R. Sun (Ed.), *The Cambridge handbook of computational psychology* (pp. 23–58). New York, NY: Cambridge University Press.
Wickelgren, W. A. (1966). Phonemic similarity and interference in short-term memory for single letters. *Journal of Experimental Psychology, 71,* 396–404.

Received June 1, 2007
Revision received February 2, 2009
Accepted February 4, 2009 ■

Postscript: More Problems With Botvinick and Plaut's (2006) PDP Model of Short-Term Memory

Jeffrey S. Bowers and Markus F. Damian
University of Bristol

Colin J. Davis
Royal Holloway, University of London

In our commentary we demonstrated that Botvinick and Plaut's (2006) model of immediate serial recall catastrophically fails when familiar letters are tested in untrained positions within a list (Simulation 2), and a modified version of their model with a distributed letter coding scheme also fails to recall familiar (and novel) letters when tested in untrained positions (Simulation 7). That is, short-term memory (STM) did not generalize to all possible test sequences. We argued that these failures reflect a fundamental limitation of the conjunctive coding schemes used in parallel distributed processing (PDP) models of cognition. Indeed, these constraints have inspired *symbolic* models of cognition that rely on context-independent representations of items in long-term memory (LTM; e.g., a representation for the letter *A*, unspecified by position within a list) and a dynamic (short-term) process of binding these items to a given role (e.g., a dynamic process of binding the letter *A* to a given position) in order to generalize more broadly.

Botvinick and Plaut (2006) rejected these claims and reported a simulation in which a new version of their model recalls familiar (and novel) items in novel positions. However, it is important to note the conditions in which this model succeeded. It included 30 input–output units, with the first 10 units coding for the onset, the next 10 units the vowel, and the final 10 units the coda. Each syllable was defined by activating one onset, one vowel, and one coda unit, and the model was trained on 999 out of a possible 1,000 ($10 \times 10 \times 10$) syllables. Their critical finding was that the model could recall the untrained item without difficulty (in all positions). What Botvinick and Plaut did not emphasize, however, was that the model was trained on all the letters in all positions of the list. So, in principle, the model could recall novel syllables (and familiar syllables in untrained positions) by recalling familiar phonemes in trained positions. For example, if the untrained syllable was *SAM*, then the model could recall SAM in Position 1 of a list by learning and activating the following trained conjunctive codes: *S*-onset-in-list-Position-1, *A*-vowel-in-list-Position-1, and *M*-coda-in-list-Position-1. Indeed, that is what the model has done.

To further highlight the generalization constraints associated with these learned conjunctive codes, we ran two new simulations. First, we developed a modified model in which the first 10 units were reserved for onsets, the next six for vowels, and the final 10 for codas (resulting in $10 \times 6 \times 10$ or 600 possible syllables). We trained the model for 3 million trials on lists of up to nine syllables taken from a random set of 300 syllables but excluded 32 syllables that included the phoneme *R* in the coda position (henceforth *R*-syllables; *R* represented by input and output Unit 17). We then trained the model for another 2 million trials, during which *R*-syllables were allowed to appear in Position 1 but not in other positions. This constitutes a general replication of the procedure we reported in our Simulation 7 but using a similar representational structure as Botvinick and Plaut's new simulation. At test, the model was presented with 1,000 lists of six syllables (taken from the vocabulary of 300) that all contained one random *R*-syllable in list Positions 1–6. As can be seen in Figure P1, when the model had not been trained on *R*-syllables, it catastrophically failed on these lists. During the additional training with the *R*-syllables, the model slowly developed a position-specific knowledge of these items: Performance improved for the *R*-syllables in first position, but the model continued to catastrophically fail when these syllables were presented in other positions. This pattern of performance is just as we reported in Simulation 7. More strikingly, in a second simulation, we trained the model on the 300 syllables with no restrictions except that the *R*-syllables were not permitted to occur in list Position 1. After 3 million training trials, the model could recall lists of six syllables that contained one *R*-syllable as long as it did not occur in Position 1. That is, when the model was tested on 1,000 lists of six syllables, its recall performance was 2.5%, 49.4%, 46.9%, 45.8%, 47.0%, and 45.5% when the *R*-syllable occurred in Positions 1–6, respectively. So, learning to recall *CAR* in list Positions 2–6 did not allow the model to recall *CAR* in Position 1. Botvinick and Plaut endorse our claim that anyone who can recall the sequence *ree-B* should also succeed in recalling the sequence *B-ree*. But as we have demonstrated here, PDP models do not exhibit the same position invariance.

These findings suggest that our new model succeeded (to the extent that it did) by relying on learned phoneme–position conjunctive codes. To test this more directly, we trained it on a random sample of 500 of the possible 600 syllables for 4 million trials and then tested it on 1,000 lists of syllables composed of

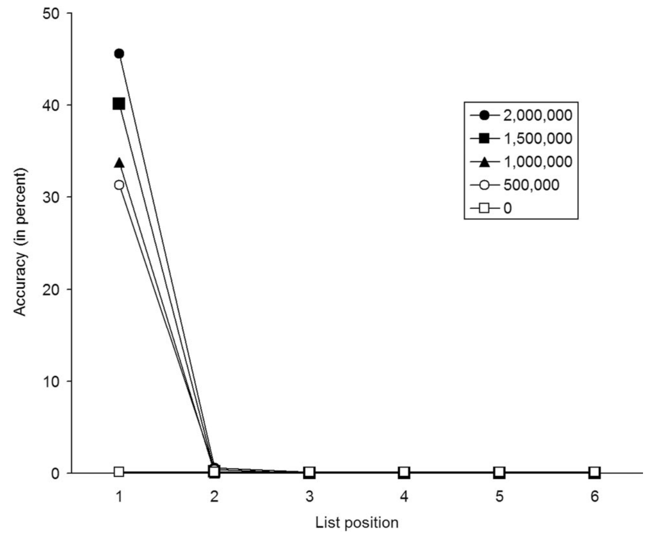

Figure P1. Performance of our modified model when it was first trained for 3 million trials on lists of syllables that excluded the phoneme *R* (*R*-syllables) and then trained another 2 million trials when the *R*-syllables were free to occur in Position 1 but not in other positions. Performance was assessed on 1,000 lists of six syllables that all contained one *R*-syllable in various positions (1–6) after various levels of training: immediately after the 3 million trials in which the *R*-syllables were untrained and following an additional 500,000, 1 million, 1.5 million, and 2 million training trials in which the *R*-syllables were free to occur in Position 1.

familiar or unfamiliar syllables that varied in length. If lists of syllables are recalled on the basis of phoneme–position conjunctive codes (e.g., the syllable *SAM* at the start of the list is coded by coactivating the long-term representations for *S*-onset-in-list-Position-1, *A*-vowel-in-list-Position-1, and *M*-coda-in-list-Position-1), then the familiarity of the syllables should be irrelevant. This is indeed the case, as depicted in Figure P2. By contrast, lexical representations play a key role in supporting human STM, as revealed by a robust advantage of words over nonwords (e.g., Jefferies, Frankish, & Lambon Ralph, 2006). Another failure of the model follows directly from this. STM is sensitive to background knowledge of sequential dependencies, and this extends to the sequential dependencies between lexical items, or newly trained syllables (e.g., Botvinick & Bylsma, 2005). Indeed, the original Botvinick and Plaut model trained on 26 letters captured these sequential effects, and this was considered a key advantage of the model compared with others. But these sequential effects are lost in the modified model given that memory performance is based on remembering sequences of phonemes. In short, when a modified Botvinick and Plaut model is trained on a larger vocabulary (e.g., 100s syllables rather than 26 letters), it suffers from both under- and overgeneralization. That is, the model cannot recall familiar (or novel) syllables that include familiar phonemes in untrained positions, but as long as this constraint is avoided (by ensuring that the training extends to all phonemes in all positions), it recalls novel syllables just as well as familiar ones and untrained sequences of syllables just as well as trained sequences.

Two additional points merit brief discussion. First, Botvinick and Plaut (2006) claimed that single-cell recording data lend support to their view that STM is mediated by context-dependent

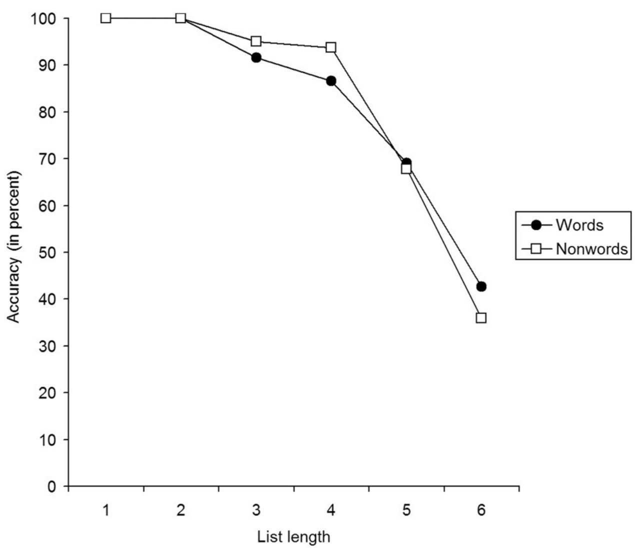

Figure P2. Performance of our modified model when it was trained for 4 million trials on random lists of syllables taken from a vocabulary of 500 of the possible 600 syllables. Performance was assessed on 1,000 lists of syllables taken from the trained (word) and untrained (nonword) sets, with list length varying from one to six syllables.

representations. But they failed to mention the evidence for context-independent representations. For example, they cited Ninokura, Mushiaske, and Tanji (2004), who reported that 30% of the relevant neurons in the lateral prefrontal cortex were selective to both position and identity (conjunctive cells). It is perhaps worth mentioning that 44% of the task-relevant neurons in this study were sensitive to list position irrespective of object identity, and 26% responded to object identity irrespective of list position (context-independent cells). Similar findings have been reported elsewhere (e.g., Averbeck, Chafee, Crowe, & Georgopoulos, 2002; Inoue & Mikami, 2006). Second, we think that Botvinick and Plaut mischaracterized Page and Norris's (1998) primacy model of STM, and they appear to have a misunderstanding regarding the representations employed in PDP and symbolic models. They claimed that the primacy model relies on conjunctive representations of items and order. But the model includes LTM representations of items that are coded independently of order, and the order of an item in a list is dynamically coded by the relative activation of the items representations. The fact that the given letter (e.g., *R*) is coded with the same unit regardless of its list position allows the model to generalize more broadly than PDP models that do rely on conjunctive representations. Indeed, all symbolic models of cognition include a process that dynamically assigns items a role, where the role could specify the position of a letter within a word (e.g., Davis, 1999), an attachment relation between object parts (e.g., Hummel & Biederman, 1992), or, in the present case, the order of items in a to-be-remembered list (e.g., Page & Norris, 1998). By contrast, Botvinick and Plaut adopted a modeling approach that binds items to roles statically, through conjunctive codes in LTM (e.g., where *R*-in-Position-1 and *R*-in-Position-2 are coded differently). By relying on a version of back-propagation, they "stipulated" that their model would learn conjunctive (context-dependent) long-term representations. The consequences are just as we predicted (see also Bowers & Davis, in press). The ball is now in their court to show that the many limitations of their model can be addressed without appealing to symbolic (context-independent) representations in LTM.

References

Averbeck, B. B., Chafee, M. V., Crowe, D. A., & Georgopoulos, A. P. (2002). Parallel processing of serial movements in prefrontal cortex. *Proceedings of the National Academy of Sciences, USA, 99,* 13172–13177.

Botvinick, M., & Bylsma, L. M. (2005). Regularization in short-term memory for serial order. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 351–358.

Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: A recurrent neural network model. *Psychological Review, 113,* 201–233.

Bowers, J. S., & Davis, C. J. (in press). Learning representations of wordforms with recurrent networks: Comment on Sibley, Kello, Plaut, and Elman (2008). *Cognitive Science.*

Davis, C. J. (1999). *The self-organising lexical acquisition and recognition (SOLAR) model of visual word recognition* (Unpublished doctoral dissertation). University of New South Wales, Sydney, Australia.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–517.

Inoue, M., & Mikami, A. (2006). Prefrontal activity during serial probe reproduction task: Encoding, mnemonic, and retrieval processes. *Journal of Neurophysiology, 95,* 1008–1041.

Jefferies, E., Frankish, C. R., & Lambon Ralph, M. A. (2006). Lexical and semantic binding in verbal short-term memory. *Journal of Memory and Language, 54,* 81–98.

Ninokura, Y., Mushiaske, H., & Tanji, J. (2004). Integration of temporal order and object information in the monkey lateral prefrontal cortex. *Journal of Neurophysiology, 91,* 555–560.

Page, M. P. A., & Norris, D. (1998). The primacy model: A new model of immediate serial recall. *Psychological Review, 105,* 761–781.