

Brian MacWhinney
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
macw@cmu.edu

Emergence from what?
Comments on Sabbagh & Gelman

Sabbagh and Gelman (S&G) present an insightful criticism of the emergentist approach to language acquisition. The analysis takes as its starting point an expressed frustration with the fact that emergentism is not packaged as a single theory or formalism. As a result, S&G decide to focus their critical attention on a particularly strong version of emergentism in which, ‘only domain-general tools are required to account for language development.’ This strong formulation of the emergentist position matches up well with the disembodied connectionism of the 1980s (Rumelhart & McClelland, 1986). However, it misrepresents the richer expressions of emergentism being developed by the authors of this volume. In particular, this ‘strong’ version fails to properly appreciate the degree to which emergentists view cognition as grounded on the body, the brain, and the social situation.

Consider a simple example from phonological development. There is a universal tendency to avoid sequences of nasal consonants followed by voiceless obstruents, as might arise in forms like ‘manpower.’ This constraint is grounded on the facts of speech production (Huffman, 1993) and figures prominently in recent elaborations of Optimality Theory (Kager, 1999). Languages use at least five phonological processes to deal with

this problem. These processes include nasal substitution, post-nasal voicing, denasalization, nasal deletion, and vowel epenthesis. Initially, children may apply a variety of these processes (Bernhardt & Stemberger, 1998). Which processes are preserved and which are dropped out will depend on the shape of the target language, whether it be Indonesian, Quechua, Toba Batak, English, or Kelantan Malay. In the terms used by S&G, each of these phonological processes is a small emergentist ‘buzzsaw’ cutting patterns that are shaped not by some innate cognitive ‘blueprint,’ but by the emergent facts of articulatory phonology.

This example of nasal assimilation constitutes a prototypical case of how emergence operates in language. However, even in this simple prototypical case, emergentism must make reference to the body. Because of the way that S&G have characterized the ‘strong’ emergentist hypothesis, even this simple example has to be excluded as an instance of emergence, since the forces driving nasal assimilation are specific to the articulatory apparatus and hence ‘domain-specific.’ Surely, something is missing in an analysis that would exclude even the most basic forms of emergence as irrelevant to the ‘strong’ emergentist hypothesis. I believe that the problem is that S&G’s strong hypothesis fails to represent the actual claims of the emergentist approach.

In particular, S&G have failed to recognize the role of the body, the brain, and the social situation in emergentist accounts. Phonology is grounded on the body. Plaut & Kello show how it is also grounded on reciprocal links between hearing and speaking. Aslin et al. show how auditory learning is grounded on statistical aspects of auditory processing. Dell & Gupta show how the shape of the lexicon is grounded on the processes of lexical retrieval. MacWhinney and Goldberg show how sentence processing is grounded on working memory, constructions, and perspective-taking, which are in turn grounded on the brain. To attempt to build an emergentist psycholinguistics that ignores the body, the brain, and the social situation would be much like attempting to build an emergentist account of honeycomb formation that ignores the honey.

S&G might argue that a ‘retreat’ from domain-generality weakens the emergentist hypothesis and makes it untestable. I would say that, by studying the interaction between domain-specific landscapes and domain-general processes, emergentism opens up paths for detailed empirical investigations. The core lesson of the last fifty years has been that both empiricism and nativism are wrong. Empiricism is wrong because it attempts to construct the mind out of nothing but domain-general ‘buzzsaws.’ Nativism is wrong because it makes untestable assumptions about genetics and unreasonable assumptions about the hard-coding of complex formal rules in neural tissue. The battles against disembodied behaviorism were fought and won in the 1950s. The battle against complex strictly-ordered rule systems was fought and won in the 1980s. We have made great progress and these issues are no longer on the table. The task facing us now is figuring out how to link domain-general processes to domain-specific landscapes.

In this sense, S&G have put their fingers on the critical issue. Unfortunately, they have mischaracterized the dominant emergentist approach to the issue. Consider the seven specific concerns they raise regarding the emergentist accounts in EL.

1. Models of the type proposed by Allen & Seidenberg rely unjustifiably on a transparent mapping between syntactic class and semantic content.
2. Connectionist models are forced to make unjustifiable *ad hoc* assumptions in order to set parameters in the model.
3. Models of the type proposed by Merriman fail to provide an adequate account of how perceptual similarity is measured.
4. In early auditory learning, domain-general mechanisms operate on domain-specific auditory proclivities.
5. Language learning accounts of the type advocated by Snow postulate too much social knowledge.
6. MacWhinney’s attempts to ground cognition on perspective are insufficiently mechanistic.

7. Connectionist accounts assume a passive child.

These are all legitimate concerns that have also been raised elsewhere. However, these concerns are not as fundamental as S&G think. If emergentism were committed to the strong S&G hypothesis, some of these criticisms would indeed be fatal. But the real issue is not how to wire up a disembodied neural network model, but rather how to link neural network mechanisms to the body, the brain, and the society. Let us look at what this means for each of these seven issues.

First, S&G complain that models of the type developed by Allen & Seidenberg rely on the transparency of the mapping between syntactic class and semantic form. This criticism is entirely fair, but the broader conclusion is incorrect. Connectionist models are often forced to accept an undeconstructed set of input features as a starting point. One first assumes a 'standard' set of syntactic, semantic, or phonological features and then focuses on examining the ways in which these standard features give rise to emergent properties. Later, one conducts a separate set of simulations in which these features are in turn deconstructed. The assumption is that, if one can get from A to B and then later from B to C, it is reasonable to think that one can get from A to C. The problem is that, without presenting a detailed model of how to get from A to B, a connectionist account seems truly disembodied. But this is not a fundamental flaw in the approach, only a problem in the technical difficulties of the implementations. Plaut & Kello show how complex connectionist models can begin to escape from this dilemma by providing linkages between perception and action in phonological development.

Building successful neural network models is not an easy task. One cannot simply vary learning rates, hidden units, and connectivity as S&G suggest and expect to obtain meaningful simulations of developmental patterns. With generative rules and production systems, one can use increasingly powerful mechanisms to account for any potential outcome. This is not true for neural network models. Instead, researchers must try to analyze problems into their components hoping to 'divide and conquer.' Connectionists

are not happy with this unnaturally analytic approach. Although the theory focuses on interactivity and connections, the full range of interactions is sometimes technically difficult to capture in a single model.

S&G worry that a domain-general theory may be unable, in principle, to provide a proper account of perceptual similarity. Here, again, the concern is entirely justified, but the conclusion is incorrect. Merriman's word-learning model is taken as an example of a theory which attempts to build up all knowledge from a domain-general basis. Both Merriman and Smith do a remarkable job illustrating how specific conceptual knowledge arises from domain-general processes such as cue validity, similarity judgment, and attentional focusing. However, neither Merriman nor Smith would want to claim that the child is a perceptual *tabula rasa*. On the contrary, like Aslin et al., they would certainly recognize the fact that domain-general learning processes interact with domain-specific perceptual abilities.

In fact, it is this particular form of dynamical interaction that captures the essence of the emergentist position. Within each perceptual system, domain-general processes of neuronal connection and refinement operate during embryogenesis and the first months of life to produce an organized perceptual landscape that is unique to each perceptual system. Motoric learning is directed toward specific limbs and muscles. When we reach the learning of the first sounds and words, these perceptual and motoric systems have developed a complex domain-specific landscape that provides rich input into general mechanisms of the type studied by Aslin et al., Merriman, Smith, and others. The case of nasal assimilation mentioned above is an example of this type. The shape of the vocal tract and the innervation of the muscles of the tongue determine the domain-specific landscape. On top of this landscape, a system of constraint-satisfaction operates to extract the proper emergent form of a grammar as a series of soft constraints. Tesar & Smolensky (1998) present one such constraint learning model, although emergentists might prefer alternative accounts that allow for richer constraint interaction.

S&G also criticize Snow's invocation of social knowledge to account for early language learning. For a fuller explication of the logic underlying this position, readers may wish to consult the argument as developed in Ninio & Snow (1988). It is important to recognize that Snow is not claiming that social knowledge is generally in advance over physical knowledge. Physical knowledge is obviously central to the development of processes such as locomotion and object permanence. However, Snow is claiming that early social knowledge is uniquely suited to supporting the acquisition of language. It supports language acquisition by helping the child both in lexical acquisition and in the learning of phrases such as 'where's the x' and 'want a y.' Moreover, fine-tuning of the feedback and input provided to the language-learning child can have a major impact on the shape of actual language usage in children growing up in different cultures and subcultures.

S&G criticize connectionists for characterizing the child as a passive learner. However, several of the chapters in *EL* make a specific attempt to characterize the child as an active learner. For example, both Snow and Smith point to ways in which the embedding of the child in a social framework can help foster the child's active internalization of social norms. Givón shows how intentional use of language eventually becomes automated into structured linguistic patterns. MacWhinney presents an even more radically activist view of the child as continually structuring language through the dynamic and willful process of perspective-taking. Paradoxically, S&G seem almost fearful of this highly activist view, worrying that the perspective-taking account cannot be easily reduced to a connectionist model. Their fears are well-grounded. Given current technology and formalization, it is not possible to reduce MacWhinney's account to a neural network. However, Bailey, Feldman, Narayanan, & Lakoff (1997) have shown how one can use control system theory to formalize the structure of embodied perspective-taking in children's early verbs. It is unfortunate that this articulate example of embodied emergentist formalization was not included in *EL*. At the same time, there

is also growing evidence that the perspective hypothesis matches up in meaningful ways with the actual architecture of the brain. It appears that the perspective hypothesis is receiving support from neurophysiology (Rizzolatti, Fadiga, Gallese, & Fogassi, 1996), even before it can be implemented in full mechanistic detail.

S&G decided to focus on their 'strong' emergentist hypothesis because they worried that any weaker version of emergentism would prove to be empirically vacuous. In fact, the majority of new empirical work in child language acquisition is devoted to testing hypotheses that fit under the general rubric of emergentism. Researchers want to understand how the shape of grammar, lexicon, phonology, and discourse emerge from real-life pressures, as well as general learning processes. The various models that have been advanced to account for these developments are rich in falsifiable empirical predictions. For S&G to suggest that emergentists have somehow decoupled their models from empirical work mischaracterizes both the articles in this volume and the field more generally.

Let us consider a few examples. Bates & Goodman argue that the strongest predictor of the growth of grammar in the second and third year is the growth of the lexicon. This important finding, if it holds up, provides direct support for Goldberg's view of the lexicon as the seat of grammatical constructions. The claimed relation could easily be falsified. If children make use of a conceptually rich proto-grammar of the type identified by Givón, for example, we would expect a much greater divergence between lexicon and grammar. Furthermore, if Elman's notion of 'starting small' is important for language learning, then children with the least morphologically articulated early vocabulary should have the most advanced syntax.

When we turn to complex syntax, Miikkulainen & Mayberry and Allen & Seidenberg make extremely clear predictions for specific reaction time patterns in on-line processing of syntactic ambiguity. Although it is true that ambiguity resolution is highly competitive, it is likely that both of these models underestimate the extent to which

processing makes some initial commitments that must be quickly reversed. It is true that Miikulainen & Mayberry, as computer scientists, have relied on others for empirical data. However, Seidenberg & MacDonald have been at the forefront of empirical work in sentence processing for years and have no shortage of empirical tests of the claims articulated in their models. In the area of pronoun resolution, MacWhinney's perspective-taking account makes strong claims that directly contravene the accepted formal linguistic theory of c-command. The empirical predictions of the perspective-taking model involve patterns of grammaticality judgments of the type customarily studied by linguists. The same model also generates a set of predictions about relative clause processing that have received strong empirical support from offline processing data (MacWhinney & Pléh, 1988), but more equivocal support from online data. The computational model of Miikulainen & Mayberry shows how the perspective-taking mechanism proposed by MacWhinney could be implemented in neural networks without having to rely on symbolic processes such as the push-down stack. In the area of word learning, the models of Merriman, Smith, and Golinkoff et al. have generated a rich outpouring of empirical studies. The models of Gupta & Dell, Plaut & Kello, and Aslin et al. are similarly committed to detailed testing through empirical data.

To summarize, we have seen that S&G correctly identify the key issue in emergentist theory as involving the application of domain-general processes. However, in their attempt to formulate an easily falsifiable version of the account, they have suggested that emergentism relies on only domain-general tools to account for language development. By incorrectly characterizing the emergentist approach in this way, they have failed to understand the extent to which emergentism seeks to link domain-general processes to a domain-specific landscape. At the same time, S&G have underestimated the extent to which emergentist theory is focused not on the global testing of emergentism, but on the specific testing of particular emergentist accounts against rich empirical data.

The papers collected in EL represent only a small fraction of the outpouring of new work in the emergentist framework. Even this small collection illustrates that this framework is not strictly controlled by a few key players, but is open to a wide variety of related applications and interpretations. Perhaps it is this distributed nature of the enterprise that provides it with such vitality, momentum, and promise. We welcome the continuation of this highly constructed critical dialog with S&G and others, since it is only through this type of interchange that new solutions to old problems in human development can emerge.

References

- Bailey, D., Feldman, J., Narayanan, S., & Lakoff, G. (1997). Modeling embodied lexical development. *Proceedings of the 19th Meeting of the Cognitive Science Society*.
- Bernhardt, B., & Stemberger, J. (1998). *Handbook of phonological development*. San Diego, CA: Academic.
- Huffman, M. K. (1993). Phonetic patterns of nasalization and implications for feature specification. In M. K. Huffman & P. A. Krakow (Eds.), *Phonetics in Phonology V*. San Diego, CA: Academic.
- Kager, R. (1999). *Optimality Theory*. New York: Cambridge University Press.
- MacWhinney, B., & Pléh, C. (1988). The processing of restrictive relative clauses in Hungarian. *Cognition*, 29, 95-141.
- Ninio, A., & Snow, C. (1988). Language acquisition through language use: The functional sources of children's early utterances. In Y. Levy, I. Schlesinger, & M. Braine (Eds.), *Categories and processes in language acquisition* (pp. 11-30). Hillsdale, NJ: Lawrence Erlbaum.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research*, 3, 131-141.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel Distributed Processing*. Cambridge, Mass.: MIT Press.
- Tesar, B., & Smolensky, P. (1998). Learnability in optimality theory. *Linguistic Inquiry*, 29, 229-268.