Developmental Science

DOI: 10.1111/j.1467-7687.2009.00826.x

SPECIAL SECTION: COMPUTATIONAL PRINCIPLES OF LANGUAGE ACQUISITION COMMENTARIES

A core principle of studying language acquisition: it's a developmental system

Larissa K. Samuelson

Department of Psychology and Delta Center, University of Iowa, Iowa City, USA

This is an exciting time to be studying language development, as is well illustrated in the collection of papers for this special section on core computational principles of language acquisition. These papers highlight innovative approaches that examine the nitty-gritty details of process. In so doing, they reveal a set of fundamental learning processes, but they also emphasize that the subject of our focus is not just a learner and something to be learned. Rather, we seek to understand a complex system in which the information provided is uniquely tuned to the intended recipient, who is particularly ready to receive just the bits of information that will carry the system forward and prepare it for the next moment of learning. In particular, these four papers examine critical aspects of the language acquisition system – the input (Hollich and Prince), the nature of the learner (McMurray, Aslin and Toscano), the fit between the learner and input (Chemla, Mintz, Bernal and Christophe), and how learning at one level influences learning at the next (Christiansen, Onnis and Hockema).

This collection of papers highlights how the complexities of the language input and aspects of the language learner fit together such that the learner can extract patterns at one level and point in time and then submit those extracted patterns to the same types of computations at the next time point to learn even more. Thus, we can see how a simple learner in a structured, dynamic, emergent, and ever-evolving context can arrive at something amazing – sophisticated, complex language in 2–3 years! Although these papers focus on learning, they also make clear that the history of the system up to the moment of learning makes this more than just simple tabula rasa information extraction. These papers show how the complexities of learning – the flow of information through the system, the fit between the learner and the input, and the nonlinearities that the learner imposes on the input – conspire to create a process that is much more than learning. This is development (cf. Spencer, Blumberg, McMurray, Robinson, Samuelson & Tomblin, in press). Recognition of the potential of this kind of developmental explanation for language acquisition comes from an openness to the possibility that emergent, nonlinear change can arise from fundamental, basic, unsophisticated principles that combine to form a developmental cascade. This is the promise of this perspective – that greater understanding of one of the most compelling changes in early cognitive development might be understood through the combined power of computational analysis and developmental process.

Hollich and Prince (this issue) focus on the visual input that often accompanies the speech stream. They examine infants' performance in an audiovisual integration task and ask how much of infant looking behaviour in tasks that require this integration can be directly accounted for by a signal-level analysis of the stimuli. Looking behaviour is a common measure of early language understanding, but there are complicated issues surrounding how an infant's looking at (and listening to) a stimulus is related to processing of that stimulus, and how this information processing drives subsequent looking behaviour. For example, what causes a release in fixation of one stimulus and a shift to examine another – an endogenous process driven by the amount of information accumulated, an exogenous process driven by salience, or, more likely, both (see Goldberg & Schöner, 2007; Schöner & Thelen, 2006; Perone & Spencer, 2008)?

Hollich and Prince begin to examine these issues by comparing the performance of infants with that of models that vary in the particulars of the information they use and the processing they complete on the available information. They find that models that use only signal-level information about visual motion account for many of the details of infant performance. They suggest, therefore, that information contained in the signal presented to infants may serve as a source for audiovisual speech integration by infants. This line of research fits with a number of other recent lines of theory and experimentation that are taking a closer look at the relation between

Address for correspondence: Larissa K. Samuelson, Department of Psychology and Delta Center, E11 Seashore Hall, University of Iowa, Iowa City, IA 52240, USA; e-mail: larissa-samuelson@uiowa.edu

Invited commentary for B. McMurray and G. Hollich (Eds.) 'Core computational principles of language acquisition: Can statistical learning do the job?' Special section in *Developmental Science*.

the visual behaviours so commonly used as measures of infant knowledge and what infants actually process and represent in such tasks (see, for example, McMurray & Aslin, 2004; Perone, Spencer & Schöner, 2007; Schoner & Thelen, 2006). However, Hollich and Prince's focus on the fit between the audio and visual information presented to children is unique and noteworthy. Thus, although they acknowledge that this work is only a first step towards a full understanding of the processes that support infant behaviour and does not yet tell us what infants are extracting from the input, further work along these lines will undoubtedly confirm that the combined audiovisual signal is a rich source of information that children can and do use to great advantage.

The paper by McMurray, Aslin and Toscano (this issue) examines properties of the learner. Specifically they take a critical look at the sufficiency of statistical learning as an explanation of infants' acquisition of phonetic categories. Using a mixture of Gaussians implementation, they find that, in fact, statistical learning processes alone are insufficient. Competition is needed to ensure the correct partitioning of the input space. McMurray *et al.* then use a full model with competitive processes to explore issues related to the sparseness of the category mapping. This yields four core observations. First, the mechanisms that underlie the observed developmental trajectory are continuious. Second, simply counting the frequency of occurrence of different categorizations is insufficient to produce the categorical distinctions that infants make. Third, the process of forming phonetic categories is one of grouping areas of auditory space together, rather than of learning particular phonological boundaries. And finally, discrimination between two inputs can be based on a categorization of one input and the lack of a category for another. Thus, behaviour that has been previously taken to indicate that infants 'have' certain phonetic categories may not actually be based on represented categories in the learner.

By asking critical questions about the nature of the learning processes, McMurray *et al.* have provided insights to the means by which a key early language-related behaviour is realized. Moreover, they have provided an important new perspective on the result of that process. Rather than carving up the phonemic space into fixed categories into which further information will have to be fitted, infants are grouping parts of it together, and leaving other parts uncategorized. This has clear implications as the system moves forward and builds on these distinctions because it is left in a more flexible state ready to incorporate future input.

The exquisite fit between what the learner needs to acquire and what she has available to do the learning is well illustrated in the work of Chemla, Mintz, Bernal and Christophe (this issue). These authors suggest how infants could extract higher-order categorical information via simple computations on lower-level structure. They propose that the lexical category of target words can be determined by tracking the co-occurrence of two context words that surround the target. Their first experiment extends this 'frequent-frames' approach, which was previously developed on a corpus of English infant-directed speech, to French. Chemla *et al.* show that, although French may appear to be problematic for the frequent-fames approach, it is, in fact, completely amenable. A second experiment shows that not just any pair of co-occurring words near a target can be used for categorization. Rather, frames need to surround target words. A third experiment demonstrates that frames are item-specific; recursion of frames does not produce even better categorization.

Together, Chemla *et al.*'s experiments show that if an infant was tracking information about individual words, the structure of the input could produce an anchor by which that child could extract higher-level information about the words, without having to shift to explicit computations at that higher level. Moreover, this work shows that not only do we not need to build in knowledge of lexical categories prior to learning, it is actually *better for infants to start off without categories*. Infants can track information about something we know they have access to early – individual words (see Christiansen *et al.*, this issue, for example) – and get to something more – categorical structure. In this way, then, the abilities of the learner are well suited to the information provided.

Christiansen, Onnis and Hockema (this issue) provide another way in which infants could extract higher-order categorical information via simple computations on lower-level structure, and show how this process can go one step further. These authors provide a computational analysis of English child-directed speech that progresses over two stages. First, transition probabilities between phonemes were used to find words in a stream of unsegmented speech. Second, distributional information about initial and final phonemes was used to predict the syntactic category of words isolated in the first stage. Christiansen *et al.* suggest that a core computational principle of language acquisition is that the same source of information can be used to learn about multiple aspects of language structure – words *and* lexical categories. Moreover, this work elegantly illustrates how a very simple computational process – tracking patterns in the stream of continuous sounds – can move the learner to a higher-level understanding of structure, without having to build the resultant categories into the system. Thus, by performing the same simple computations on the successively more structured product of its own computations, even a relatively unsophisticated learner can capture what has previously been seen as unobtainable abstract structure.

As is clear in these papers, the field of language development has come to a greater appreciation of infants' and toddlers' amazing language *learning* abilities. This, combined with the increasing use of innovative and quite sophisticated modelling techniques, has driven a shift in focus away from descriptions of language behaviours and abilities at different developmental stages, and towards the details of developmental processes. Increasingly, the central question is what are the mechanisms of change. And in answering this question there is a corresponding move to focusing on the nature, structure and richness of the input to learners at multiple timescales and very fine-grained levels of detail. This is complemented by a greater understanding of the learner as situated in a rich and supportive context that evolves over time. Thus, when the field previously asked 'what information do language learners have to work with?' and 'what can they *learn* with that information?' the answers were 'too little' and 'not much'. More recent work such as that presented in this collection, however, suggests that the answers to these questions are actually, 'a lot' and 'quite a bit'. An important question for future work, then, will be what are the processes by which this evolution from learning to development occurs. That is, how do individual instances of in-the-moment learning coalesce and build on a longer timescale into developmental change (see also McMurray, Toscano, Horst & Samuelson, in press; Samuelson & Horst, 2008). Understanding the core computational principles of language development is an important step in this direction.

References

Chemla, E., Mintz, T.H., Bernal, S., & Christophe, A. (this issue). Categorizing words using 'frequent fames': what cross-linguistic analyses reveal about distributional acquisition strategies. *Developmental Science*, **12** (3), 397–407.

Christiansen, M.H., Onnis, L., & Hockema, S.A. (this issue). The secret is in the sound: from unsegmented speech to lexical categories. *Developmental Science*, **12** (3), 388–396.

Goldberg, J., & Schöner, G. (2007). Understanding the distribution of infant attention: a dynamical systems approach. In D.S. McNamara & J.G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 1043–1048). Austin, TX: Cognitive Science Society.

Hollich, G., & Prince, C.G. (2009). Comparing infants' preference for correlated audiovisual speech with signal-level computational models. *Developmental Science*, **12** (3), 379–387.

McMurray, B., & Aslin, R.N. (2004). Anticipatory eye movements reveal infants' auditory and visual categories. *Infancy*, **6** (2), 203–229.

McMurray, B., Aslin, R.N., & Toscano, J.C. (this issue). Statistical learning of phonetic categories: insights from a computational approach. *Developmental Science*, **12** (3), 369–378.

McMurray B.A., Horst, J.S., Toscano J., & Samuelson, L.K. (in press). Connectionist learning and dynamic processing: symbiotic developmental mechanisms. In J. Spencer, M. Thomas and J. McClelland (Eds.), *Toward a new grand theory of development? Connectionism and dynamic systems theory reconsidered*. Oxford: Oxford University Press.

Perone, S., & Spencer, J.P. (2008). Linking individual differences in looking behavior to learning and recognition. Poster presented at the *Sixteenth Biennial International Conference on Infant Studies*, Vancouver.

Perone, S., Spencer, J.P., & Schöner, G. (2007). A dynamic field theory of visual recognition in infant looking tasks. In D.S. McNamara & J.G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 1391–1396). Austin, TX: Cognitive Science Society.

Samuelson, L.K., & Horst, J.S. (2008). Confronting complexity: insights from the details of behavior over multiple timescales. *Developmental Science*, **11**, 209–215.

Schöner, G., & Thelen, E. (2006). Using dynamic field theory to rethink infant habituation. *Psychological Review*, **113** (2), 273–299.

Spencer, J.P., Blumberg, M.S., McMurray, B., Robinson, S.R., Samuelson, L.K., & Tomblin, J.B. (2009) Short arms and talking eggs: the inconvenience of understanding process. *Child Development Perspectives* (in press).

DOI: 10.1111/j.1467-7687.2009.00827.x

The learner as statistician: three principles of computational success in language acquisition

Melanie Soderstrom,[1] Erin Conwell,[2] Naomi Feldman[3] and James Morgan[3]

1. *Department of Psychology, University of Manitoba, Winnipeg, Canada*
2. *Department of Psychology, Harvard University, Cambridge, USA*
3. *Department of Cognitive and Linguistic Sciences, Brown University, Providence, USA*

Statistical learning is the new paradigm of language acquisition. A perusal of recent conference programs or journal contents reveals much work advocating – or criticizing – statistical learning. Language acquisition

Address for correspondence: Melanie Soderstrom, Department of Psychology, P404 Duff Roblin Building, 190 Dysart Road, University of Manitoba, Winnipeg, MB, Canada R3T 2N2; e-mail: M_Soderstrom@umanitoba.ca

Invited commentary for B. McMurray and G. Hollich (Eds.) 'Core computational principles of language acquisition: Can statistical learning do the job?' Special section in *Developmental Science*.