A gradually developing lexicon leads to robust emergence of phonological features in a neural network

[authors, affiliation]
[email addresses]

We present a neural network in which phonological features emerge as a result of auditory and lexical acquisition. Our model requires at least three levels of representation: an auditory-phonetic form [AudF], a phonological surface form /SF/, and a lexical form. These levels are implemented as separate layers in the neural network, which is shown in its initial state in Fig. 1:

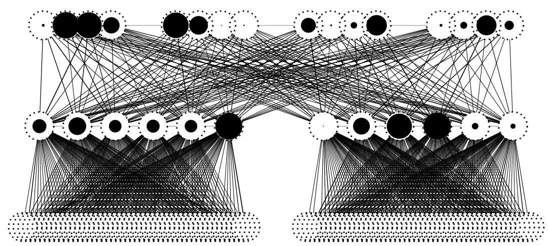

Figure 1. *The neural network before learning begins*.

At this point, the nodes in the network have small, random activities, and the SF nodes are connected to the other layers by connections with small, random weights. (More active nodes are drawn with larger black circles inside them; heavier connections are drawn with thicker lines.) All SF nodes that represent an auditory continuum are connected to all other SF nodes within the same layer by inhibitory connections, which ensure that as a node becomes active, it simultaneously deactivates the other nodes in that layer. This way, nodes become dedicated to parts of the auditory continuum ("competitive learning": Grossberg 1976; Rumelhart & Zipser 1985). The input language contains four lexical categories (the four groups of nodes in the top layer in Fig. 1), each of which can be described in terms of two auditory continua. If, for instance, these continua are VOT and sibilant centre of gravity, our input language could be the sibilant inventory [s z ʃ ʒ].

Earlier work within this model showed that categorical behavior emerges at SF after distributional learning of auditory tokens alone (Boersma, Benders & Seinhorst, 2018). In Chládková's (2014: ch. 5) simulations, the lexicon was in place from the onset of learning; in [[placeholder for self-reference]] the network abruptly went from the prelexical stage to the lexical stage halfway through the learning process. However, neither scenario is realistic, as perceptual learning and lexical acquisition in infants appear to proceed in tandem (a.o. Swingley 2008). For this reason, we created a neural network in which lexical information becomes available to the learner in a gradual fashion from the onset of learning: in this implementation, the activations of the nodes in the lexicon layer increase throughout the learning process following a logistic function. The 8,000th learning step is shown in Figure 2a (next page). At this point, the top-down influence from the lexicon is not particularly strong yet: the activation of the nodes at the lexicon layer is at about 50% of its maximum value. We can already see that the nodes at SF are becoming dedicated to certain parts of the auditory continua: coincidentally, the leftmost nodes at both SF layers are the only nodes to have become connected to the very leftmost parts of the corresponding auditory continua. Figure 2b shows the 16,000th and last learning step of the same network, when the lexicon is fully in place.

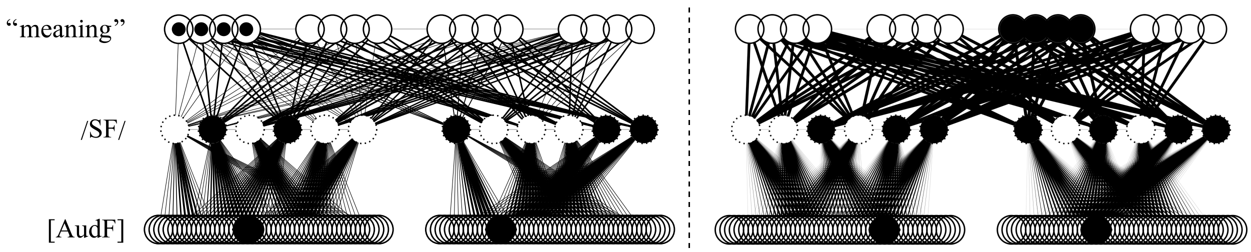

Figure 2. (a) *The 8,000th learning step.* (b) *The 16,000th learning step.*

When learning is complete (Fig. 2b), we can see that the nodes at SF have become specialized: for instance, on the left continuum, SF nodes 1, 2, and 4 have become specialized in the left region of the corresponding AudF, and nodes 3, 5 and 6 in the right region. Fig. 2b also shows that the SF nodes that were recruited by a region of AudF are connected to those AudF nodes in the exact same way, which was not yet the case in Fig. 2a: this happened as a result of the top-down information from the lexicon. The specialization of the nodes at SF indicates that categorical behavior has emerged there, since only two possible activation patterns exist: for instance, at the left SF, there is one pattern in which nodes 1, 2 and 4 are on while nodes 3, 5 and 6 are off; and one in which nodes 3, 5 and 6 are on while nodes 1, 2, 4 are off. Whereas these features had to be inferred in earlier computational models of auditory distributional learning (Feldman et al. 2009; McMurray et al. 2009), they are directly observable in our model, which allows us to track their development from being more audition-based (as in Fig. 2a) to being audition- and lexicon-based (as in Fig. 2b).

The logistic curve in our simulations is perhaps not entirely appropriate, as lexical acquisition in the language-learning infant does not necessarily proceed smoothly. However, the emergence of features in our model is robust to fluctuations: categorical behavior still emerges at SF if we vary the steepness and midpoint of the logistic curve, or if we add random noise to the activations at the lexical layer. Features also emerge irrespective of the shape of the auditory input distributions: for instance, if the pooled auditory distribution only has one peak, the top-down information from the lexicon ensures the emergence of distinct categories (cf. Ter Schure, Junge & Boersma 2016), while distributional learning alone would fail to create a contrast at SF.

References:

Boersma, Paul, Titia Benders & Klaas Seinhorst (2018). *Neural network models for phonology and phonetics*. Manuscript, University of Amsterdam.

Chládková, Kateřina (2014). *Finding phonological features in perception*. Dissertation, University of Amsterdam.

Feldman, Naomi, Thomas Griffiths & James Morgan (2009). Learning phonetic categories by learning a lexicon. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, 2208–2213.

Grossberg, Stephen (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics* 23: 121–134.

McMurray, Bob, Richard Aslin & Joe Toscano (2009). Statistical learning of phonetic categories: insights from a computational approach. *Developmental Science* 12: 369–378.

Rumelhart, David & David Zipser (1985). Feature discovery by competitive learning. *Cognitive Science* 9: 75–112.

Swingley, Daniel (2008). The roots of the early vocabulary in infants' learning from speech. *Current Directions in Psychological Science* 17 (5): 308–312.

ter Schure, Sophie, Caroline Junge & Paul Boersma (2016). Semantics guide infants' vowel learning: computational and experimental evidence. *Infant Behavior and Development* 43: 44–57.