

How do children learn words they have never been taught?

Juan Segurola · All Rights Reserved

QUE · Article 038 · Simple Questions, Deep Answers · Version 0.1 EN

Children learn many words without someone first arranging a lesson and providing a definition. That does not mean they learn without evidence. Speech, recurring situations and interactions supply information even when nobody explicitly announces what is being taught. The scientific problem is to explain how a learner extracts useful structure from input that is often ambiguous.

There are at least two problems hidden in the title. First, a child must find candidate words in a stream of sound. Second, the child must connect a candidate with something it can mean. Recognising a recurring sound sequence does not by itself identify its referent, and identifying an object does not reveal which stretch of speech names it.

1. Speech does not arrive with printed spaces

Written sentences make word boundaries look obvious. Spoken sequences do not simply provide an audible blank at every place where a reader would insert a space. One possible cue is statistical: some sound sequences recur together more consistently than sequences crossing boundaries.

Saffran, Aslin and Newport's 1996 experiment showed that eight-month-old infants could discriminate sequences after exposure to an artificial speech stream in which statistical relationships between neighbouring syllables provided segmentation information. [1] The result concerns a controlled segmentation task. It does not demonstrate that infants understood the artificial items as names for objects, or that distributional information alone explains an entire language.

An invented example illustrates the cue. Suppose one syllable is followed by a second in 30 of its 30 occurrences, while another pairing occurs in 10 of 30. The conditional probabilities are 1 and one-third. A learner sensitive to the difference has evidence for a grouping, although the grouping is not automatically a word. Real input contains repeated phrases, variable pronunciations and other regularities that complicate that inference.

2. A name and a scene still leave a puzzle

Suppose an adult says “dax” while two unfamiliar objects are visible. Even if the child has segmented the sound correctly, the intended object remains uncertain. If the objects always appear together, repeated co-occurrence by itself will not distinguish them. More exposure helps only when it supplies discriminating information.

Smith and Yu tested twelve- and fourteen-month-old infants with individually ambiguous word-scene pairings. Their study found learning of word-referent mappings across those pairings, supporting the possibility that information can accumulate across situations rather than being resolved completely at each encounter. [2]

The key phrase is “across situations”. Learning need not wait for a single perfectly unambiguous teaching episode. But the experiment does not imply unlimited memory, perfect tracking of every association or a mechanism indifferent to what the infant notices. Statistical information must be available to the learner, not merely present in a researcher's complete record.

3. A small case of accumulating evidence

Consider a simplified learner that hears the same label in three scenes. The candidate objects are A and B, then A and C, then A and D. Under the explicit assumptions that the label denotes exactly one stable object and that its referent appears in every scene, the intersection of the three candidate sets is A. No single scene settled the mapping; the combination does.

This is a logical illustration, not the algorithm established by the infant study. Relax an assumption and its answer can change. If an adult sometimes refers to an absent object, simple intersection can fail. If the label describes a property shared by several objects, looking only for one object gives the wrong search space. If A and B always occur together, the observations do not decide between them.

Those limitations are scientifically productive. They show what a model has to add to handle natural situations: a representation of possible meanings, a way to deal with missing or misleading input, and some account of what is retained. A successful toy learner establishes possibility under assumptions, not a complete developmental explanation.

4. Interaction provides more than a word count

In an ordinary interaction, an adult may point, orient toward an object, repeat a phrase or respond to the child's action. These events can make some candidate interpretations more relevant than others. They also change the information available for statistical learning. Statistical structure and interaction need not be competing explanations if interaction helps organise the observations from which learning proceeds.

That proposal should still be tested rather than assumed. For example, one could compare the same word-object exposures with different attentional cues, measure where infants look, and ask whether later learning follows exposure count or attended exposure. A design of that kind separates two explanations that would look identical if only the final vocabulary score were recorded.

The studies cited here demonstrate bounded capacities in controlled conditions. They do not decide how much of language is learned through each route across languages, ages or communicative environments. They also do not settle the broad nature-versus-nurture debate. Showing sensitivity to input does not explain the developmental capacities that make that sensitivity possible.

5. What learning a word ultimately requires

Choosing an object when a label is spoken is one outcome. Extending a word appropriately to a new example, understanding it in a sentence and using it to communicate are additional achievements. An experiment should be described by the outcome it actually measures. Calling every successful choice full word knowledge inflates a useful result.

The answer to the title is that explicit lessons are only one source of evidence. Children can extract regularities in speech and across ambiguous situations, while interactions can help organise what they encounter. What remains under investigation is how these resources combine over development. The conceptual mistake is to infer that no formal teaching means no information, then replace a learnable problem with a mystery.

References

[1] Saffran, J. R., Aslin, R. N., and Newport, E. L. Statistical learning by 8-month-old infants. *Science* 274(5294), 1926–1928 (1996). <https://doi.org/10.1126/science.274.5294.1926>

[2] Smith, L., and Yu, C. Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition* 106(3), 1558-1568 (2008). <https://doi.org/10.1016/j.cognition.2007.06.010>

Related reading

Article 037: Why can an hour feel short or long?