Induction in Language Learning

Induction is a type of reasoning that involves deriving general principles from specific examples. It contrasts with deduction, in which specific examples are derived from generalities. Deductions can be made with certainty: from a general statement like *All men are mortal* one can conclude that a specific man, Bob, is mortal. Induction, however, is probabilistic: there are always many logically possible ways to **generalize** on the basis of specific facts. Nearly all aspects of language must be learned by induction: children hear specific instances of phonemes, morphology, words, and sentences, from which they must derive the general principles of their language.

<u>Induction in word learning</u>. The fact that induction is inherently logically under-constrained is especially evident in word learning, as was pointed out by the philosopher W.V.O. Quine. He postulated the existence of an explorer visiting a new tribe. Upon seeing a rabbit, one of the natives exclaims "Gavagai!" One might conclude from this that *gavagai* means "rabbit", but it could also mean any number of other possibilities, including "Look!", "Let's go hunting!", or even "An undetached rabbit part." In fact, there are an infinite number of possibilities consistent with the meaning of any word. Children are not led astray by all of these logical possibilities: at most, a child might entertain a few hypotheses ("rabbit", "tail", or "dinner"). The fact that people's generalizations are constrained in this way is an indication that children come to the language learning task with certain **biases**. The open questions in language development are what biases people have for learning different aspects of language, and where those biases come from.

Word learning is an inductive problem because children hear individual instances of a word (like *cat*), and they must then decide what other items should also be called *cat.* One theory for how children decide this is that they are guided by certain principles. For instance, the **whole object bias** suggests that children usually presume that words apply to entire objects rather than their parts. Thus, the word *cat* probably picks out the whole animal rather than just the cat's tail or whiskers. Another suggested principle is the **mutual exclusivity bias**, which guides induction in the presence of known words. Suppose a child is shown two items, one which she knows the label for ("ball") and one which she doesn't. If she hears a new label ("cat") she will guess that it applies to the item whose label she does not already know. Although this bias is sometimes false (cats may also be called *pets* or *animals*), it usefully guides children's inductions about word meaning.

One possible source of word learning biases is that they may be derived from general rules of pragmatic inference. For example, a child might arrive at the mutual exclusivity bias by reasoning about the intended communicative goals of the person they are talking to: "I know the round thing is called a *ball*, so if mom had intended me to get that she would have said *ball*. Since she didn't, she probably meant the furry meowing thing when she said *cat*." A related type of reasoning can guide inductions based on multiple examples. If a child is shown three different Dalmatians and told they are all called a *fep*, he is likely to conclude that *fep* means "Dalmation" rather than "dog"; if it meant "dog" then presumably some of the given examples would have been a different kind of dog, like a chihuahua or a poodle.

Word learning biases may also be acquired from the structure of the input. By two years old, English-speaking children have a **shape bias**; they assume that nouns are organized by shape rather than color or texture. Shown a star-shaped, speckled toy called a *dax*, a child will presume that the star-shaped matte toy (but not the conical speckled one) is also called a *dax*. This assumption parallels the structure of their vocabulary, since count nouns in English tend to be classified by shape.

Another constraint on induction in word learning comes about because most words are labeled in many different situations. In the *gavagai* example, if the explorer hears the word again in a situation involving a deer, it becomes less likely that *gavagai* means "rabbit" and more likely that it means "Let's go hunting!" Indeed, both children and adults are able to use cross-situational information like this to constrain their inductions about word meaning.

Induction in grammar learning. Induction also plays a key role in forming inferences about grammar. One early argument demonstrating that humans have grammar-learning biases was presented by E. Mark Gold, who showed mathematically that even a perfect learner who heard sentences from a language could only ever be 100% certain about the grammar if the language was finite. The reason is that if the language allows an infinite number of grammatical sentences (as all languages do) it is impossible to be completely sure which sentences that have not yet been heard are grammatical. One possibility, called the **subset principle**, is for a learner to select a grammar that can produce *only* sentences she has heard; but if she does this, she might miss some sentences that she hasn't heard but are allowed. It is clear that children do not follow the subset principle: they generalize beyond their input, producing sentences like "he goed to school" that they almost certainly did not hear. Conversely, if a learner selects a grammar that produces some sentences that she hasn't heard, she might be wrong about them and never converge on the correct grammar.

This observation has spurred a great deal of research into whether children receive or make use of **negative evidence** -- information about what constructions are *not* grammatical. Such evidence would overcome Gold's problem, since children could use it to eliminate grammars that produce sentences that they haven't heard. However, negative evidence is rare, and when it does exist it is often statistical and noisy or mostly ignored by children. It is therefore usually thought that children's grammatical inductions are guided (at least in large part) by learning biases instead.

Different biases may guide different aspects of grammar learning. One method for identifying what biases exist is known as the **argument from the poverty of the stimulus**. The idea is that if children appear to make a kind of generalization, but haven't heard constructions that would support such a generalization, then they must have a bias telling them to make it anyway. This argument has been applied to many areas of grammar learning, including forming certain kinds of questions in English as well as correctly using the word "one" in constructions like "he liked the red one." Arguments from the poverty of the stimulus can be useful ways of identifying biases, but are often contentious. This is because it can be difficult to identify what kind of input (or how much) is necessary to support any given generalization, and therefore to determine when it is sufficiently lacking to require the presence of a learning bias instead.

Grammatical induction has been studied widely in the realm of verb learning, since verbs only occur in limited constructions. For instance, the verb *put* must be associated with three phrases: the actor doing the putting (*the boy*), the thing that is put (*the ball*) and the place it is put (*on the chair)*. Learning which verbs can occur in which constructions involves inductive generalization, since people presented with novel verbs will spontaneously use them in constructions they haven't heard them in before. This sort of generalization involves some mixture of verb-specific learning combined with construction-specific learning, and may be guided by principles about the language as well as the meaning of the verbs in question.

Modeling induction in language learning. Computational models of language learning can be useful for explaining how induction occurs. Any model must balance between precisely matching the data and generalizing beyond it. Some models, like **connectionist networks**, achieve this generalization due to the structure of the network. Other approaches, like **Bayesian models**, describe an "optimal" way of generalizing based on the nature of the problem and the goals of the learner. These models tend to favor hypotheses that fit the data, but do so simply, without having to postulate complicated representations. Human performance is then compared to model predictions to determine to what extent people behave like those learners. Both Bayesian and connectionist models have been applied successfully to the induction problems discussed above.

Amy Perfors
University of Adelaide

See also: Bayesian Inference in Word Learning; Cross-Situational Observation in Word Learning; Over-Extension and Under-Extension in Word Learning; Over-Generalization of Grammatical Constructions; Poverty of the Stimulus Argument; Shape Bias in Word Learning; Word Learning: Constraints

Further Readings

Bloom, Paul. *How children learn the meanings of words*. MIT Press. (2000)

Gold, E. Mark. "Language identification in the limit." *Information and Control,* v.10(5): 447-474. (1967)

Goldberg, Adele. *A construction grammar approach to argument structure*. University of Chicago Press. (1995)

Pinker, Steven. *Learnability and cognition: The acquisition of argument structure*. MIT Press. (1989)

Quine, Willard V.O. *Word and object.* MIT Press. (1960)