


Provided for non-commercial research and educational use.
Not for reproduction, distribution or commercial use.

This article was originally published in the Encyclopedia of Infant and Early Childhood Development, 2nd edition, published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use, including without limitation, use in instruction at your institution, sending it to specific colleagues who you know, and providing a copy to your institution's administrator.



All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<https://www.elsevier.com/about/policies/copyright/permissions>

From Goldin-Meadow, S., 2020. Theories of Language Acquisition. In: Benson, J.B. (Ed.), Encyclopedia of Infant and Early Childhood Development, 2nd edition. vol. 3, Elsevier, pp. 356–364. <https://dx.doi.org/10.1016/B978-0-12-809324-5.23585-2>
ISBN: 9780128165126

Copyright © 2020 Elsevier Inc. unless otherwise stated. All rights reserved
Elsevier

Theories of Language Acquisition[☆]

Susan Goldin-Meadow, University of Chicago, Departments of Psychology and Comparative Human Development, Chicago, IL, United States

© 2020 Elsevier Inc. All rights reserved.

Theoretical Accounts of Language-Learning	356
Behaviorist Accounts	356
Nativist Accounts	357
Social/Cognitive Accounts	358
Connectionist Accounts	358
Constrained Learning	359
Constrained Invention	360
Is Language Innate?	362
Innateness Defined as Genetic Encoding	362
Innateness Defined as Developmental Resilience	362
Language Is Not a Unitary Phenomenon	363
References	363

The simplest technique to study the process of language-learning is to do nothing more than watch and listen as children talk. In the earliest studies, researcher parents made diaries of their own child's utterances (e.g., [Stern and Stern, 1907](#); [Leopold, 1939–1949](#)). The diarist's goal was to write down all of the new utterances that the child produced. Diary studies were later replaced by audio and video samples of talk from a number of children, usually over a period of years. The most famous of these modern studies is [Roger Brown's \(1973\)](#) longitudinal recordings of Adam, Eve, and Sarah.

Because transcribing and analyzing child talk is so labor-intensive, each individual language acquisition study typically focuses on a small number of children, often interacting with their primary caregiver at home. However, advances in computer technology have made it possible for researchers to share their transcripts of child talk via the computerized Child Language Data Exchange System (CHILDES, <https://childes.talkbank.org>). Because this system makes available many transcripts collected by different researchers, a single researcher can now call upon data collected from spontaneous interactions in naturally occurring situations across a wide range of languages, and thus test the robustness of descriptions based on a small sample. In addition, naturalistic observations of children's talk can always be, and often are, supplemented with experimental probes that are used with larger numbers of subjects.

Thus, it is possible, although time-consuming, to describe what children do when they acquire language. The harder task is to figure out how they do it.

Many theories have been offered to explain how children go about the process of language-learning. This chapter begins by reviewing the major accounts. We will find that, although there is disagreement among the theories in the details, all modern day accounts accept the fact that children come to the language-learning situation prepared to learn. The disagreement lies in what each theory takes the child to be prepared with. Do children come to language-learning equipped with a general outline of what language is? Or do they come instead with a set of processes that will lead to the acquisition of language (and language alone)? Or do they come with a set of processes that will lead to the acquisition of any skill, including language? The chapter then describes recent theoretical and experimental approaches that have been and are currently being applied to the problem of determining the constraints that children bring to language-learning, and ends with an analysis of what it might mean to say that language is innate.

Theoretical Accounts of Language-Learning

Behaviorist Accounts

Consistent with the psychological theories of that era, prior to the late 1950s language was considered just another behavior that can be acquired by the general laws of behavior, such as associative learning, reinforcement, and imitation. Consider, for example, associative learning, a general learning process in which a new response becomes associated with a particular stimulus. Association seems like a natural way to explain how children learn the words of their language – children learn labels by associating a heard

[☆]*Change History:* March 2019. Updating additions made to "Social/Cognitive Accounts" and "Connectionist Accounts.", Shortened "Is Language Innate?" and "Innateness Defined as Developmental Resilience," and eliminated "Mechanisms that Could Lead to Resilience.", References added. Research described in this chapter was supported by grant no. R01 DC00491 from the National Institute on Deafness and Other Communication Disorders and grant no. BCS 1654154 from the National Science Foundation.

This article is an update of S. Goldin-Meadow, Language Acquisition Theories, Editor(s): Marshall M. Haith, Janette B. Benson, Encyclopedia of Infant and Early Childhood Development, Academic Press, 2008, Pages 177-187.

word with a visible object. But it's not so simple. Quine's (1960) famous theoretical puzzle highlights the problem: Imagine that you are a stranger in a foreign country with no knowledge of the local language. A native says "gavagai" while pointing at a rabbit running in the distance. You try to associate the new response "gavagai" with a particular stimulus, but which stimulus should you choose? The entire rabbit? Its tail? Its ears? The running event? The possibilities are limitless and associative learning solves only a small piece of the problem.

Imitation and reinforcement were also proposed as mechanisms by which children could learn the grammatical "habits" that comprise language. However, even the most cursory look at how children learn language reveals that neither of these mechanisms is sufficient to bring about language-learning. Children learn the language to which they are exposed and, in this broad sense, learn language by imitation. But do children model the sentences they produce after the sentences they hear? Some do, but many children are not imitators (Bloom et al., 1974). Moreover, the children who are imitators do not learn language any more quickly than the non-imitators. Even the children who routinely imitate do not copy everything they hear – they are selective, imitating only the parts of the sentences that they are able to process at that moment. Thus, imitation is guided as much by the child as by the sentences the child hears.

What about the responses of others to children's sentences? Do children learn to use sentence types that parents reinforce? Parents might positively reinforce sentences their children produce that are grammatically correct and negatively reinforce sentences that are grammatically incorrect. In this way, the child might be encouraged to produce correct sentences and discouraged from producing incorrect ones. There are two problems with this account. The first is that parents do not typically respond to their children's sentences as a function of the grammatical correctness of those sentences (Brown and Hanlon, 1970). Parents tend to respond to the content rather than the form of their children's sentences. Second, even if children's grammatically correct sentences were treated differently from their grammatically incorrect sentences, it is still up to the child to determine what makes the correct sentences correct. For example, if the child says the grammatically correct sentence, "I colored the wall blue," and mother responds with positive reinforcement (thus ignoring the sentence's troubling content and focusing on its form), the child still has to figure out how to generalize from the sentence; she needs to understand the patterns that generate the sentence in order to recognize that one analogous sentence ("I saw the wall blue") is not grammatically correct while another ("I pounded the clay flat") is. In other words, there would still be a great deal of inductive work to be done even if children were provided with a set of correct sentences from which to generalize.

The behaviorist account of language was dealt a devastating blow with the publication in 1959 of Noam Chomsky's review of B.F. Skinner's *Verbal Behavior*. Chomsky argued that adult language use cannot be adequately described in terms of sequences of behaviors or responses. A system of abstract rules underlies each individual's knowledge and use of language, and it is these rules that children acquire when they learn language. When viewed in this way, the language acquisition problem requires an entirely different sort of solution.

Nativist Accounts

The premise of the Chomskian perspective is that children are learning a linguistic system governed by subtle and abstract principles without explicit instruction and, indeed, without enough information from the input to support induction of these particular principles (as opposed to other principles) – Plato's problem or the poverty of the stimulus argument. Chomsky went on to claim that if there is not enough information in the input to explain how children learn language, the process must be supported by innate syntactic knowledge and language-specific learning procedures (Chomsky, 1965). The theory of Universal Grammar (UG) formulates this *a priori* knowledge in terms of principles and parameters that determine the set of possible human languages. UG is assumed to be part of the innately endowed knowledge of humans. The principles of UG provide a framework for properties of language, often leaving several (constrained) options open to be decided by the data the child comes in contact with. For example, word order freedom is a parameter of variation. Some languages (English) mandate strict word orders; others (Russian, Japanese) list a small set of admissible orders; still others (Warlpiri, an Australian aboriginal language) allow almost total scrambling of word order within a clause. Input from a given language is needed for learners to set the parameters of that language.

One important aspect of this theory is that setting a single parameter can cause a cluster of superficially unrelated grammatical properties to appear in the language. For example, the *null-subject parameter* involves a number of properties: whether overt subjects are required in all declarative sentences (*yes* in English, *no* in Italian), whether expletive elements such as "it" in "it seems" or "there" in "there is" are exhibited (*yes* in English, *no* in Italian), whether free inversion of subjects is allowed in simple sentences (*no* in English, *yes* in Italian), and so on. The prediction is that the input necessary to set the null-subject parameter results in the simultaneous alignment of all of these aspects within a child's grammar (Hyams, 1989). There is, at present, controversy over whether predictions of this sort are supported by the child language data. Moreover, in linguistic theory, the search for universal principles and parameters has been augmented by the view that innate linguistic knowledge is as economic a system as possible (the Minimalist Program, Chomsky, 1993).

Innate knowledge of the principles underlying language, however, is not sufficient to account for how children acquire language. How are children to know what a noun or a subject is in the specific language they are learning? Children obviously need to identify subjects and verbs in their language before they can determine whether the two are strictly ordered in that language, and before they can engage whatever innate knowledge they might have about how language is structured. Thus, in addition to innate syntactic knowledge, children also need learning procedures, which may themselves be language-specific.

One example of a possible learning procedure is a set of rules linking semantic and syntactic categories (Pinker, 1989). Under this hypothesis, children are assumed to know innately that agents are likely to be subjects, that objects affected by action are likely to be direct objects, and so on. All they need do is identify (using context) the agent in a scene; the linking rules allow them to infer

that the term used to refer to that agent is the subject of the sentence. Their innate knowledge about how these elements are allowed to be structured can then take over. Again, controversies exist over whether child language data support these assumptions (e.g., ergative languages do not straightforwardly link agents with subjects and yet are easily acquired by young children, e.g., Ochs, 1982; see also Bowerman, 1990).

Social/Cognitive Accounts

The nativist position entails essentially two claims: (1) at least some of the principles of organization underlying language are language-specific and not shared with other cognitive systems, and (2) the procedures that guide the implementation of these principles are themselves innate, that is, centered in the child and not the child's environment. Note that, while these two claims often go hand-in-hand, they need not. One can imagine that the principles underlying linguistic knowledge might be specific to language and, at the same time, implemented through general, all-purpose learning mechanisms (although such mechanisms must be more complex than the mechanisms behaviorist accounts have offered). This position has come to be known as a social or cognitive account of language-learning (for a recent example, see Bohn et al., 2018).

For example, by observing others' actions—where they look, how they stand, how they move their hands and faces—we can often guess their intentions. Young children could use this information to help them narrow down their hypotheses about what a word means. In fact, if a speaker looks at an object while uttering a novel word, a child will assume that the speaker's word refers to that object, even if the child herself is not looking at the object (Baldwin, 1993). In other words, children can use general cues to speaker intent to guide their guesses about language.

Children can also use the language itself as a cue to meaning. If the syntactic structures of the sentences that children hear reflect, at least to some extent, their meanings (as Gleitman and her colleagues have found, e.g., Landau and Gleitman, 1985; Fisher et al., 1991), then children could use that structure to bootstrap their way into meaning, a process known "syntactic bootstrapping". There is now quite a bit of evidence that even young children can break into language via syntactic bootstrapping (e.g., Naigles et al., 1993).

Children do not sound like adults when they begin to speak – there clearly is developmental work that needs to be done. The question is what type of work is required? One possibility, favored by some nativists, is that children have in place all of the grammatical categories and syntactic principles they need; they just lack the operating systems that will allow those principles to run. The developmental work to be done does not, under this view, involve a changing grammatical system.

Another view suggests that the child's language changes dramatically during development, transforming from a system based on semantic categories to one based on syntactic categories. This transformation could be determined maturationally or guided by innate linking rules. However, the transformation could also result from an inductive leap children make on the basis of the linguistic data available to them, in conjunction with the cognitive and/or social skills they bring to the task – this inductive leap is at the heart of all social or cognitive accounts of language acquisition.

Cognitive underpinnings are obviously necessary but they may not be sufficient for the onset of linguistic skills. For example, the onset of gesture + speech combinations that convey two elements of a proposition ("open" + point at box) precedes the onset of two-word combinations ("open box") by several months, suggesting that the cognitive ability to express two semantic elements is not the final stumbling block to two-word combinations (Iverson and Goldin-Meadow, 2005). More than likely, it is the difficulty of extracting linguistic patterns from the input that presents the largest problem.

Social and cognitive accounts claim that there is enough information in the linguistic input children hear, particularly in the context of the supportive social environments in which they live, to induce a grammatical system. Ample research indicates that adults alter the speech they direct to their children. Speech to children (often called *motherese* or *child-directed speech*) is slower, shorter, higher-pitched, more exaggerated in intonation, more grammatically well formed, and more directed in content to the present situation than speech addressed to adults (Snow, 1972). And children pay particular attention to this fine-tuned input, interpreting it in terms of their own biases or operating principles (e.g., paying attention to the ends of words, Newport et al., 1977).

However, one problem that arises with postulating motherese as an engine of child language-learning is that child-directed speech may not be universal. In many cultures, children participate in communicative interactions as overhearers (rather than as addressees) and the speech they hear is not likely to be simplified in the same ways. Nevertheless, children in these cultures become competent users of their grammatical systems in roughly comparable time frames (Ochs and Schieffelin, 1995). These observations suggest that there may be many developmental routes to the same end – a reasonable conjecture given the robustness of language.

One very interesting possibility that skirts the problem that children do not universally receive simplified input is that the children may do the simplifying themselves. For example, young children's memory limitations may make them less able to recall entire strings of words or all the sounds that make up the individual words. As a result, they do the analytic work required to abstract linguistic regularities on a smaller, filtered data base (the "less is more" hypothesis, Newport, 1990; see also Elman, 1993). This filtering may be just what children require to arrive at their linguistic systems. Moreover, it is a general process that children around the globe presumably bring, in equal measure, to the language-learning situation.

Connectionist Accounts

Connectionism is a movement in cognitive science whose goal is to explain human intellectual abilities using artificial neural networks (also known as neural nets). Neural networks are simplified models of the brain composed of large numbers of units

(the analogs of neurons) and weights that measure the strength of the connections between those units. In a connectionist account, behavior is shaped by selective reinforcement of the network of interconnected units. Under this view, language development is a process of continuously adjusting the relative strengths of the connections in the network until linguistic output resembles linguistic input (Elman, 2001).

In a sense, connectionism is more of a technique for exploring language-learning than an explanatory account. But connectionism does involve some theoretical assumptions. For example, most connectionist models are based on the assumption that language (like all other cognitive skills) can be explained without recourse to rules.

Connectionism offers a tool for examining the tradeoff between the three components central to all theories of language learning – environment (input to the system), structures the child brings to the learning situation (architectures of the artificial system), and learning mechanisms (learning algorithms). For example, a great deal of linguistic structure is assumed to be innate on the nativist account. Connectionism can provide a way to explore how much structure needs to be built in to achieve learning, given a particular set of inputs to the system and a particular set of learning mechanisms. As another example, networks have examined how variations in structure—for example the size of the memory span—influence language learning. Networks arrive at appropriate generalizations from strings of sentences *only if* the memory span of the network for previously processed words begins small and gradually increases (reminiscent of the “less is more” hypothesis described earlier, see Elman, 1993). In principle, connectionism is agnostic on the question of whether the architecture of the system (the child) or the input to the system (the environment) determines the relative strengths of each connection. However, in practice, most connectionists emphasize the importance of input. And the unanswered question is what determines the units that are to be connected in the first place.

Machine learning is another approach that can be applied to language learning (e.g., Freudenthal and Alishahi, 2014). Machine learning is a field of Artificial Intelligence whose goal is to create software applications that, over time, become more accurate on a complex task and, in this sense, learn the task. By simulating the process of child language learning, computational models have the potential to reveal which linguistic representations are learnable from the input children actually receive. To the extent that a computational model yields the acquisition patterns that children actually exhibit, that model provides insight into plausible mechanisms underlying human language development. Although many current day models are now successful in learning the patterns to which they are exposed, they often have trouble going beyond those patterns and fail to make the generalizations that a human learner would (see Lake et al., 2017, for a review). The problem may grow out of the desire to put as little structure into the model as possible; that is, to build a simple architecture that has few predispositions. But we know that children come to language learning with at least some predispositions, and our computational models of language learning might fare better if we tried to build evidence-based predispositions into them.

Constrained Learning

Another approach to understanding how human children are prepared to learn language is to assume they come with general biases to process information in a particular way. This view suggests that children's strong inclination to structure communication in language-like patterns results from their general processing biases coming into contact with natural language input.

The language that children learn must, at some level, be inferable from the data that are out there. After all, if linguists manage to use language data to figure out the grammar of a language, why can't children? But linguists can be selective in ways that children can not. Linguists do not have to weigh all pieces of data equally; they can ask informants what an utterance means and whether it is said correctly. Linguists also have at their disposal a great deal of data at one time. The question is – what kinds of learning mechanisms can we realistically impute to children that will allow them to make sense of the data they receive as input?

One learning mechanism that has been proposed is known as statistical learning. The assumption underlying this mechanism is that children are sensitive to the patterns in their input and can perform rapid and complex computations of the co-occurrences among neighboring elements in that input. By performing statistical computations over a corpus, children can pick out the recurring patterns in the data and thus are less likely to be misled by individual counter-examples.

However, children must also face the problem that a corpus can be analyzed in many different ways. How do they know which computations to perform on a given corpus? Perhaps children are only able to perform a limited set of computations. If so, this limitation would effectively narrow down the range of possible patterns that could be extracted from a database. Thus, one way that children may be prepared to learn language is that they come to language-learning ready to perform certain types of computations and not others.

To discover which computations young language-learning children are able to perform, we can provide them with a corpus of data constructed to exhibit a pattern that can be discovered using a particular computation. If the children then extract the pattern from the data, we know that they are able to perform this type of computation on a corpus. As an example, 8-month old infants were exposed to a corpus of nonsense words playing continuously on an audiotape for 2 min (Saffran et al., 1996). The corpus was arranged so that the transitional probabilities between sounds were 1.0 inside words, but 0.33 across words—that is, a particular syllable followed another specific syllable 100% of the time within words, and a particular syllable followed another specific syllable 33% of the time between words. The only way the infant could figure out what the words in the corpus were was to (i) pay attention to these transitional probabilities and (ii) assume that sequences with high probabilities are likely to be inside words and that sequences with low probabilities are likely to be the accidental juxtapositions of sounds at word boundaries. The infants not only listened differentially to words versus non-words, but they were able to discriminate between words and

part-words (part-words contained the final syllable of one word and the first two syllables of another word; they were thus part of the corpus the infants heard but had different transitional probabilities from the words). The 8-month-olds were not merely noting whether a syllable sequence occurred – they were inducing a pattern from the sounds they had heard and using a mechanism that calculates statistical frequencies from input to do so.

Infants are thus sensitive to the transitional probabilities between sounds and can use them to segment speech into word-like units. Can this simple mechanism be used as an entry point into higher levels of linguistic structure? If, for example, children can use transitional probabilities between words (or word classes) to segment sentences into phrases, they could then use this phrasal information as a wedge into the syntax of their language. In other words, children may be able to go a long way toward inducing the structure of the language they are learning by applying a simple procedure (tabulating statistical frequencies) to the data that they receive. A related domain-general approach that has been taken to the problem is the Bayesian inference framework, a tool for combining prior knowledge (probabilistic versions of constraints) and observational data (statistical information in the input) in a rational inference process (e.g., Xu and Tenenbaum, 2007). The theoretical assumption underlying all of these approaches is that children come to language-learning equipped with processing strategies that allow them to induce patterns from the data to which they are exposed.

The open question is – how sophisticated do the data-processing strategies have to be in order for children to induce the patterns of their language from the input that they actually receive? Can children get by with the ability to calculate transitional probabilities, building up larger and larger units over developmental time? Or are there units over which children are more, or less, likely to calculate transitional probabilities? For example, children may (or may not) be able to calculate statistical probabilities over units that are not immediately adjacent (i.e., dependencies between units that are at a distance from one another, e.g., in the sentence, “the cats on the couch are beautiful,” the verb *are* is plural because it depends on *cats*, the subject of the sentence, which occurs several words earlier). Some of the constraints that children exhibit during language-learning may come from the processing mechanisms they bring to the situation.

Two questions are frequently asked about language processing mechanisms: (1) Are the mechanisms that children apply to language-learning task specific and unique to language, or are they used in other domains as well? (2) Are the mechanisms children apply to language-learning species-specific and unique to humans, or are they used by other species as well?

The task-specificity question can be addressed with respect to statistical learning by providing children with non-language input that is patterned (e.g., musical patterns, visual patterns) and observing whether young infants can discover those patterns. They do (e.g., Slone and Johnson, 2015) – suggesting that calculating transitional probabilities is a general purpose data-processing mechanism that children apply to their worlds. The species-specificity question can be addressed with respect to statistical learning by exposing non-humans to the same type of language input that the human infants heard, and observing whether they can discover the patterns. It turns out that cotton-top tamarin monkeys can extract word-like units from a stream of speech sounds just as human infants do (Hauser et al., 2001). But, of course, tamarin monkeys do not acquire human language. The interesting question, then, is where do the monkeys fall off? What types of computations are impossible for them to perform? This theoretically motivated paradigm thus allows us to determine how the mechanisms children bring to language constrain what they learn, and whether those constraints are specific to language and specific to humans.

Constrained Invention

When children apply their data-processing mechanisms to linguistic input, the product of their learning is language. But what if a child was not exposed to linguistic input? Would such a child be able to invent a communication system and, if so, would that communication system resemble language? If children are able to invent a communication system that is structured in language-like ways, we must then ask whether the constraints that guide language-learning are the same as the constraints that guide language-invention.

Language was clearly invented at some point in the past and was then transmitted from generation to generation. Was it a one-time invention, requiring just the right assembly of factors, or is language so central to being human that it can be invented anew by each generation? This is a question that seems impossible to answer – today's children do not typically have the opportunity to invent a language, as they are all exposed from birth (and perhaps even before birth since babies can perceive some sounds in utero) to the language of their community. The only way to address the question is to find children who have *not* been exposed to a human language.

It turns out that there are children who are unable to take advantage of the language to which they are exposed. These children are congenitally deaf with hearing losses so severe that they cannot acquire the spoken language that surrounds them, even with intensive instruction. Moreover, they are born to hearing parents who do not know a sign language and have not placed their children in a situation where they would be exposed to one. These children lack an accessible model for human language. Do they invent one?

The short answer is yes (Feldman et al., 1978). These children are able to communicate with the hearing individuals in their worlds, and use gesture to do so. This is hardly surprising since all hearing speakers gesture when they talk. The surprising result is that the deaf children's gestures do not look like the gestures that their hearing parents produce. The children's gestures have language-like structure; the parents' gestures do not (Goldin-Meadow, 2003).

The children's gestures are composed of parts, akin to morphemes in conventional sign languages (as in spoken language, morphemes in sign language are the smallest unit that conveys meaning). For example, an *eat* gesture (an O-shaped hand moved

back and forth at the mouth) is composed of a handshape morpheme (the O-shape denoting grasping a small object) and a motion morpheme (the back-and-forth movement denoting jabbing at the mouth). The children combine these gestures into sentence-like strings that are structured with grammatical rules for deletion and order. For example, to ask an adult to share a snack, one child pointed at the snack, gestured *eat*, and then pointed at the adult. He typically placed gestures for the object of an action before gestures for the action, and gestures for the agent of an action after. Importantly, the children's gesture systems are generative (Goldin-Meadow and Yang, 2017). For example, one child concatenated gestures conveying several propositions about snow shovels within the bounds of a single gesture sentence: that they are used to dig, that they are used when boots are worn, that they are used outside and kept downstairs. The gesture systems have parts of speech (nouns, verbs, adjectives), are used to make generic statements (as in the snow shovel example), and are used to tell stories about the past, present, future and the hypothetical. The children even use their gestures to talk to themselves and to talk about their own gestures.

In contrast, the children's hearing parents use their gestures as all speakers do. Their sloppily formed gestures are synchronized with speech and are rarely combined with one another. The gestures speakers produce are meaningful, but they convey their meanings holistically, with no componential parts and no hierarchical structure (Goldin-Meadow et al., 1995).

The theoretically interesting finding is not that the deaf children communicate with their gestures, but that their gestures are structured in language-like ways. Indeed, the children's gestures are sufficiently language-like that they have been called *homesigns* (Goldin-Meadow, 2003). It is important to note that the deaf children could have used mime to communicate – for example, miming eating a snack to invite the adult to join the activity. But they do not. Instead, they produce discrete, well-formed gestures that look more like beads on a string than a continuous unsegmentable movement.

Segmentation and combination are at the heart of human language, and they form the foundation of the deaf children's gesture systems. But segmentation and combination are not found in the gestural input children receive from their hearing parents. Thus, the deaf children could not easily have taken data-processing strategies of the sort that have been hypothesized and applied them to the gestural input they receive in order to arrive at their homesign systems. Although it is clear that children who are exposed to a language must be applying data-processing strategies to the sentences they hear in that language in order to acquire it (e.g., Lieven, 2016), it is equally clear that children can arrive at a language-like system through other routes. A communication system structured in language-like ways seems to be overdetermined in humans.

The deaf children invented the rudiments of language without a model to guide them. But they did not invent a full-blown linguistic system – perhaps for good reason. Their parents wanted them to learn to talk and thus did not share the children's gesture systems with them. As a result, the children's systems were one-sided – they produced language-like gestures to their parents, but received non-linguistic co-speech gestures in return.

What would happen if such a deaf child were given a partner with whom to develop language? Just such a situation arose in the 1980's in Nicaragua when deaf children were brought together in a group for the very first time. The deaf children had been born to hearing parents and, like the deaf children described earlier, presumably had invented gesture systems in their individual homes. When they were brought together, they needed to develop a common sign language, which has come to be called *Nicaraguan Sign Language* (NSL) (Kegl et al., 1999). The distance between the homesigns invented by individual children without a partner and the sign system created by this first cohort of NSL can tell us which linguistic properties require a shared community in order to be introduced into human language.

But Nicaraguan Sign Language has not stopped growing. Every year, a new cohort of deaf children enters the group and learns to sign among their peers. Each new cohort of signers has as its input the sign system developed by the previous cohort. And each new cohort continues to adapt the system so that the product becomes even more language-like. The properties of language that crop up in this second and subsequent cohorts are properties that depend on passing the system through fresh minds – linguistic properties that must be transmitted from one generation to the next in order to be introduced into human language (e.g., Senghas and Coppola, 2001). NSL is not unique among sign languages – it is likely that all sign languages came about through a similar process (see, for example, Sandler et al., 2005).

Because sign languages are processed by eye and hand rather than ear and mouth, we might have expected them to be structured differently from spoken languages. But they are not. Sign languages all over the world are characterized by the same hierarchy of linguistic structures (syntax, morphology, phonology) and thus draw on the same human abilities as spoken languages. Moreover, children exposed to sign language from birth acquire that language as naturally as hearing children acquire the spoken language to which they are exposed, achieving major milestones at approximately the same ages (Newport and Meier, 1985). However, the manual modality makes sign languages unique in at least one respect. It is easy to use the manual modality to invent representational forms that can be immediately understood by naïve observers (e.g., indexical pointing gestures, iconic gestures). Thus, sign languages can be created anew by individuals and groups, and are particularly useful in allowing us to determine whether language-creation is constrained in the same ways that language-learning is.

Computational and robotic experiments offer another approach to the problem of language invention. These studies explore whether communication systems with properties akin to those found in natural human language can emerge in populations of initially language-less agents. There are two traditions in this work. The first functional approach assumes that linguistic structure arises as a solution to the problem of communication, for example, as a way of limiting search through possible interpretations (e.g., Steels, 2016). The second structural approach does not rely on communication pressure to motivate change but rather examines the emergence of structure as the system is passed from one user (or one generation of users) to the next (e.g., Kirby et al., 2008). Studies in this second tradition have found that a compositional system with recursion, grammatical categories, and word order inevitably results from passing an initially unstructured communication system through generations of learners. These are just the properties

found in the deaf children's homesign gesture systems – but the homesign systems are not passed through a series of learners and are instead invented by individual children who are the sole users of their systems. Once again, we find that there is more than one way to arrive at language-like structure. In general, modeling studies, combined with observations of actual cases of language-learning and language-invention, can help us appreciate the range of circumstances under which language-like structure can arise and the mechanisms responsible for that structure.

Is Language Innate?

Children are likely to come to language-learning constrained to process the language data they receive in certain ways and not in others. The constraints could be specifically linguistic, but they need not be. Constraints are assumed to be internal to the child at the moment when a particular skill is acquired. But are they innate?

Innateness Defined as Genetic Encoding

The problem of innateness has been addressed repeatedly and elegantly in other disciplines, especially ethology, and many definitions of innateness have been proposed (Wimsatt, 1986). One of the most common, albeit not the earliest, definitions of an innate behavior is that it have a genetic base. Some have claimed evidence for grammar genes – not for a single gene responsible for all the circuitry underlying grammar, but for a set of genes whose effects are relevant to the development of the circuits underlying parts of grammar (Pinker, 2003). The dispute is whether the genes are specific to the grammatical aspects of language.

What might it mean to claim that language has a genetic base? At one level, the claim is obviously true. All humans who are genetically intact have comparable linguistic systems – in the same way that all human bodies have two arms and two legs. The details of the arms of any two unrelated individuals (their length, width, definition) are likely to differ (and those differences may or may not be grounded at the genetic level) but the basic two-ness and structure of the arm is constant across all genetically intact humans – so too for language.

So why then (assuming we are not geneticists) should we care about the genetic base of language-learning? Perhaps we should not. Of all of the very large number of definitions and criteria that have, over the years and over the disciplines, been applied to the term innate, one could argue that the definition that is least central to the notion's core is having a genetic base. A useful definition of innate need not be anchored in genetic mechanisms.

Innateness Defined as Developmental Resilience

An alternative definition of an innate behavior is that it is developmentally resilient. A behavior is developmentally resilient if its development, if not inevitable, is overdetermined in the species; that is, it is a behavior likely to be developed by each member of the species even under widely varying circumstances. The way we traditionally explore the boundaries for the development of a behavior is to manipulate the conditions under which that behavior is typically developed, extending the range until the behavior no longer appears. For obvious ethical reasons, we cannot tamper with the circumstances under which children learn language. But we can take advantage of variations in language-learning conditions that occur naturally, and thus explore the boundary conditions under which language development is possible.

Although language-learning is not infinitely resilient (e.g., when human children are raised under inhumane circumstances, they do not develop language, Curtiss, 1977), it does display a surprising amount of robustness in the face of either external (environmental) or internal (organic) variation (Bishop and Mogford, 1988; Goldin-Meadow, 1989). As an example of environmental variation, children around the globe differ in how much, when, and what types of language they receive – not to mention the fact that, in each culture, the child is exposed to a model of a different language (Ochs and Schieffelin, 1995). Indeed, many children are exposed to input from two different languages and must learn both at the same time. Despite the broad range of inputs, children in all corners of the earth learn language and at approximately the same pace (Slobin, 1985-1997).

Language-learning is also resilient in the face of many organic variations from the norm, variations that alter the way children process whatever input they receive. For example, blind children live in a nonvisual world that is obviously different from the sighted child's world, and that offers a different spectrum of contextual cues to meaning. However, this difference has little impact on language-learning in the blind child (Landau and Gleitman, 1985). Organic variation can be much more severe and still result in relatively intact language-learning. For example, grammar-learning in the earliest stages can proceed in a relatively normal manner and at a normal rate even in the face of unilateral ischemic brain injury (Feldman, 1994).

There may be no greater testament to the resilience of language than the fact that children can invent language in the absence of a model for language. A combination of internal factors (the fact that the children are profoundly deaf and cannot acquire a spoken language) and external factors (the fact that the children have not been exposed to a conventional sign language) together create the unusual language-learning circumstances in which the deaf children described earlier find themselves. Despite their lack of a model for language, these children still communicate in language-like ways.

In sum, language development can proceed in humans over a wide range of environments and a wide range of organic states, suggesting that the process of language development may be buffered against both environmental and organic variations. No one factor seems to be ultimately responsible for the course and outcome of language development in humans, a not-so-surprising

result given the complexity and importance of human language. It looks as though there is a basic, resilient form that human communication naturally gravitates toward, and a variety of developmental paths that can be taken to arrive at that form. In this sense, language development in humans can be said to be characterized by equifinality – a term coined in the early 1900's by the embryologist Driesch (as reported in [Gottlieb, 1995](#)) to describe a process by which a system reaches the same outcome despite widely differing input conditions. No matter where you start, all roads lead to Rome.

Language Is Not a Unitary Phenomenon

Until now we have been discussing language as though it were a unitary phenomenon, as though it were obvious what the appropriate unit of analysis for language is. However, it is clear that language is not a unitary whole, particularly when it comes to issues of resilience and innateness.

Children who are not exposed to a conventional language model create communication systems that contain some, but not all, of the properties found in natural human languages ([Goldin-Meadow, 2003](#)). Thus, the absence of a conventional language model appears to affect some properties of language more than others. Even when linguistic input is present, it is more likely to affect rate of acquisition for certain properties of language than for others ([Newport et al., 1977](#)). Further, when language is acquired off-time (i.e., in late childhood or adolescence) certain properties of language are likely to be acquired and others are not ([Newport, 1990](#)). Thus, some properties of language are relatively resilient, while others are relatively fragile. Moreover, there is some evidence that the same properties of language (e.g., using the order of words to convey who does what to whom) are resilient across many different circumstances of acquisition – acquisition without a conventional language model, acquisition with varying input from a language model, and acquisition late in development after puberty. Thus, language as a whole need not be said to be innate.

The definition of innate that best fits the language-learning data is developmental resilience. This notion operationalizes innateness by specifying the range of organisms and environments in which language-learning can take place. There clearly are limits on the process of language development – children raised without human interaction do not develop language. But the process of language development can proceed in children with a range of limitations and in children raised in environments that vary radically from the typical. What we see in exploring this resilience is that certain aspects of language are central to humans – so central that their development is virtually guaranteed, not necessarily by a particular gene but by a variety of combinations of genetic and environmental factors. In this sense, language is innate.

References

- Baldwin, D.A., 1993. Infants' ability to consult the speaker for clues to word reference. *J. Child Lang.* 20 (2), 395–418.
- Bishop, D., Mogford, K. (Eds.), 1988. *Language Development in Exceptional Circumstances*. Churchill Livingstone, New York.
- Bloom, L., Hood, L., Lightbown, P., 1974. Imitation in language development: if, when, and why. *Cogn. Psychol.* 6 (3), 380–420.
- Bohn, M., Zimmermann, L., Call, J., Tomasello, M., 2018. The social-cognitive basis of infants' reference to absent entities. *Cognition* 177, 41–48.
- Bowerman, M., 1990. Mapping thematic roles onto syntactic functions: are children helped by innate linking rules? *Linguistics* 28, 1253–1289.
- Brown, R., 1973. *A First Language*. Harvard University Press, Cambridge, MA.
- Brown, R., Hanlon, C., 1970. Derivational complexity and order of acquisition in child speech. In: Hayes, J.R. (Ed.), *Cognition and the Development of Language*. Wiley, N.Y.
- Chomsky, N., 1959. A review of B.F. Skinner's *Verbal Behavior*. *Language* 35 (1), 26–58.
- Chomsky, N., 1965. *Aspects of the Theory of Syntax*. MIT Press, Cambridge, MA.
- Chomsky, N., 1993. A Minimalist Program for Linguistic Theory. MIT Occasional Papers in Linguistics No. 1. MIT Working Papers in Linguistics, Cambridge, MA.
- Curtiss, S., 1977. Genie: A Psycholinguistic Study of a Modern-Day 'wild Child'. Academic Press, N.Y.
- Elman, J.L., 1993. Learning and development in neural networks: the importance of starting small. *Cognition* 48 (1), 71–99.
- Elman, J.L., 2001. Connectionism and language acquisition. In: Tomasello, M., Bates, E. (Eds.), *Essential Readings in Developmental Psychology*. Language Development: The Essential Readings. Blackwell Publishing, Malden, pp. 295–306.
- Feldman, H.M., 1994. Language development after early unilateral brain injury: a replication study. In: Tager-Flusberg, H. (Ed.), *Constraints on Language Acquisition Studies of Atypical Children*. Erlbaum Associates, N.J., pp. 75–90.
- Feldman, H., Goldin-Meadow, S., Gleitman, L., 1978. Beyond Herodotus: the creation of a language by linguistically deprived deaf children. In: Lock, A. (Ed.), *Action, Symbol, and Gesture: The Emergence of Language*. Academic Press, N.Y.
- Fisher, C., Gleitman, L., Gleitman, H., 1991. On the semantic content of subcategorization frames. *Cogn. Psychol.* 23, 331–392.
- Freudenthal, D., Alishahi, A., 2014. Computational models of language development. In: Brooks, P., Kempe, V. (Eds.), *Encyclopedia of Language Development*. Sage Publications.
- Goldin-Meadow, S., 1989. Review of *Language Development in Exceptional Circumstances*. *Lang. Speech* 32 (3), 285–289.
- Goldin-Meadow, S., 2003. *The Resilience of Language*. Psychology Press, N.Y.
- Goldin-Meadow, S., Mylander, C., Butcher, C., 1995. The resilience of combinatorial structure at the word level: morphology in self-styled gesture systems. *Cognition* 56, 195–262.
- Goldin-Meadow, S., Yang, C., 2017. Statistical evidence that a child can create a combinatorial linguistic system without external linguistic input: implications for language evolution. *Neurosci. Biobehav. Rev.* 81 (Part B), 150–157.
- Gottlieb, G., 1995. A systems view of psychobiological development. In: Magnusson, D. (Ed.), *Individual Development over the Lifespan: Biological and Psychosocial Perspectives*. Cambridge University Press, N.Y.
- Hauser, M.D., Newport, E.L., Aslin, R.N., 2001. Segmentation of the speech stream in a non-human primate: statistical learning in cotton-top tamarins. *Cognition* 78, B53–B64.
- Hyams, N., 1989. The null subject parameter in language acquisition. In: Jaeggli, O.A., Safir, K.J. (Eds.), *The Null Subject Parameter*, *Studies in Natural Language and Linguistic Theory*, vol. 15. Springer, Dordrecht.
- Iverson, J.M., Goldin-Meadow, S., 2005. Gesture paves the way for language development. *Psychol. Sci.* 16, 368–371.
- Keigl, J., Senghas, A., Coppola, M., 1999. Creation through contact: sign language emergence and sign language change in Nicaragua. In: DeGraff, M. (Ed.), *Language Creation and Language Change: Creolization, Diachrony, and Development*. MIT, Cambridge, MA, pp. 179–237.
- Kirby, S., Cornish, H., Smith, K., 2008. Cumulative Cultural Evolution in the Laboratory: an experimental approach to the origins of structure in human language. *Proc. Natl. Acad. Sci. Unit. States Am.* 105 (31), 10681–10686.

- Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S., 2017. Building machines that learn and think like people. *Behav. Brain Sci.* <https://doi.org/10.1017/S0140525X16001837>.
- Landau, B., Gleitman, L.R., 1985. *Language and Experience: Evidence from the Blind Child*. Harvard University Press, Cambridge, MA.
- Leopold, W.F., 1939–1949-.). In: *Speech Development of a Bilingual Child: A Linguist's Record*, vols. 1–4. Northwestern University Press, IL.
- Lieven, E., 2016. Usage-based approaches to language development: where do we go from here? *Lang. Cognit.* 8 (3), 346–368.
- Naigles, L., Gleitman, H., Gleitman, L., 1993. Syntactic bootstrapping in verb acquisition. In: Dromi, E. (Ed.), *Language and Cognition: A Developmental Perspective*. Ablex, N.J.
- Newport, E.L., 1990. Maturation constraints on language learning. *Cogn. Sci.* 14, 11–28.
- Newport, E.L., Gleitman, H., Gleitman, L.R., 1977. Mother, I'd rather do it myself: some effects and non-effects of maternal speech style. In: Snow, C.E., Ferguson, C.A. (Eds.), *Talking to Children: Language Input and Acquisition*. Cambridge University Press, Cambridge.
- Newport, E.L., Meier, R., 1985. Acquisition of American sign language. In: Slobin (Ed.), *The Cross-linguistic Study of Language Acquisition*. Erlbaum, N.J.
- Ochs, E., 1982. Ergativity and word order in Samoan child language. *Language* 58, 646–671.
- Ochs, E., Schieffelin, B., 1995. The impact of language socialization on grammatical development. In: Fletcher, P., MacWhinney, B. (Eds.), *The Handbook of Child Language*. Blackwell Publishers, Oxford.
- Pinker, S., 1989. *Learnability and Cognition. The Acquisition of Argument Structure*. MIT Press, Cambridge, MA.
- Pinker, S., 2003. *The Language Instinct: How the Mind Creates Language*. Penguin, UK.
- Quine, W.V., 1960. *Word and Object*. MIT Press, Cambridge, MA.
- Saffran, J.R., Aslin, R.N., Newport, E.L., 1996. Statistical learning by 8-month-old infants. *Science* 274 (5294), 1926–1928.
- Sandler, W., Meir, I., Padden, C., Aronoff, M., 2005. The emergence of grammar: systematic structure in a new language. *Proc. Natl. Acad. Sci. Unit. States Am.* 102, 2661–2665.
- Senghas, A., Coppola, M., 2001. Children creating language: how Nicaraguan Sign Language acquired a spatial grammar. *Psychol. Sci.* 12, 323–328.
- Slobin, D.I. (Ed.), 1985–1997. *A Cross-Linguistic Study of Language Acquisition*, vols. 1–5. Erlbaum Associates, N.J.
- Slone, L.K., Johnson, S.P., 2015. Infants' statistical learning: 2- and 5-month-olds' segmentation of continuous visual sequences. *J. Exp. Child Psychol.* 133, 47–56.
- Snow, C.E., 1972. Mothers' speech to children learning language. *Child Dev.* 43, 549–565.
- Steels, L., 2016. Agent-based models for the emergence and evolution of grammar. *Phil. Trans. R. Soc. B* 371. <https://doi.org/10.1098/rstb.2015.0447>.
- Stern, C., Stern, W., 1907. *Die Kindersprache*. Verlag von Johann Ambrosius Barth, Leipzig.
- Wimsatt, W., 1986. Developmental constraints, generative entrenchment, and the innate-acquired distinction. In: Bechtel, W. (Ed.), *Integrating Scientific Disciplines*. Martinus Nijhoff, Dordrecht, pp. 185–208.
- Xu, F., Tenenbaum, J.B., 2007. Word learning as Bayesian inference. *Psychol. Rev.* 114, 245–272.