

Susan Carey

*Bootstrapping &
the origin of concepts*

All animals learn. But only human beings create scientific theories, mathematics, literature, moral systems, and complex technology. And only humans have the capacity to acquire such culturally constructed knowledge in the normal course of immersion in the adult world.

There are many reasons for the differences between the minds of humans and other animals. We have bigger brains, and hence more powerful information processors; sometimes differences in the power of a processor can create what look like qualitative differences in kind. And of course human beings also have language – the main medium for the cultural transmission of acquired knowledge. Comparative studies of humans and other primates suggest that we differ from them as well in our substantive cognitive abilities – for example, our capacity for causal anal-

ysis and our capacity to reason about the mental states of others. Each of these factors doubtless contributes to our prodigious ability to learn.

But in my view another factor is even more important: our uniquely human ability to ‘bootstrap.’ Many psychologists, historians, and philosophers of science have appealed to the metaphor of bootstrapping in order to explain learning of a particularly difficult sort – those cases in which the endpoint of the process transcends in some qualitative way the starting point. The choice of metaphor may seem puzzling – it is self-evidently impossible to pull oneself up by one’s own bootstrap. After all, the process I describe below is not impossible, but I keep the term because of its historical credentials and because it seeks to explain cases of learning that many have argued *are* impossible.

Sometimes learning requires the creation of new representational resources that are more powerful than those present at the outset. Early in the cultural history of mathematics, for instance, the concept of the number included only positive integers: with subsequent development the concept came to encompass zero, rational numbers (fractions), negative numbers, irrational numbers like pi, and so on.

Susan Carey, professor of psychology at Harvard University, has played a leading role in transforming our understanding of cognitive development. A Fellow of the American Academy since 2001, she is the author of numerous articles and essays and the book “Conceptual Change in Childhood” (1985).

© 2004 by the American Academy of Arts & Sciences

Bootstrapping is the process that underlies the creation of such new concepts, and thus it is part of the answer to the question: What is the origin of concepts?

Individual concepts are the units of thought. They are constituents of larger mental structures – of beliefs that are formed out of them and of systems of representation such as intuitive theories. Concepts are individuated on the basis of two kinds of considerations: their reference to different entities in the world and their role in distinct mental systems of inferential relations.

How do human beings acquire concepts? Logic dictates three parts to any explanation of the origin of concepts. First, we must specify the innate representations that provide the building blocks of the target concepts of interest. Second, we must describe how the target concepts differ from these innate representations – that is, we must describe developmental change. And third, we must characterize the learning mechanisms that enable the construction of new concepts out of the prior representations.

Claims about all three parts of the explanation of the origin of concepts are highly controversial. Many believe that innate representations are either perceptual or sensory, while others (including myself) hold that humans and other animals are endowed with some innate representations with rich conceptual content. Some researchers also debate the existence, even the possibility, of qualitative changes to the child's initial representations. One argument for the impossibility of such radical changes in the course of development is the putative lack of learning mechanisms that could explain them. This is the gap that my appeal to bootstrapping is meant to fill.

To make clear both what the problem is, and what role bootstrapping may play

in solving it, I will examine how children acquire one specific set of concepts: the positive integers – i.e., concepts such as *one, two, three, nine, eighteen*, etc.


Before they acquire language, infants form several different types of representation with numerical content, at least two of which they share with other vertebrate animals.

One, described by Stanislas Dehaene in his delightful book *The Number Sense*, uses mental symbols that are neural magnitudes linearly related to the number of individuals in a set. Because the symbols get bigger as the represented entity gets bigger, they are called analog magnitudes. Figure 1 gives an external analog magnitude representation of number, where the symbol is a line, and length is the magnitude linearly related to number. Mental computations using these symbols include comparison, to establish numerical difference or equality, and also addition and subtraction.

Mental analog magnitudes represent many dimensions of experience – for example, brightness, loudness, and temporal duration. In each case as the physical magnitudes get bigger, it becomes increasingly harder to discriminate between pairs of values that are separated by the same absolute difference. You can see in figure 1 that it is harder to tell that the symbol for seven is different from (and smaller than) that for eight than it is to tell that the symbol for two is different from (and smaller than) that for three. Analog magnitude representations follow Weber's law, according to which the discriminability of two values is a function of their ratio.

You can confirm for yourself that you have an analog magnitude system of representation of number that conforms to Weber's law. Tap out as fast as you can without counting (you can prevent your-

Figure 1
Analog magnitude models
Number represented by a quantity linearly related
to the cardinal value of the set



self from counting by thinking ‘the’ with each tap) the following numbers of taps: 4, 15, 7, and 28. If you carried this out several times, you’d find the mean number of taps to be 4, 15, 7, and 28, with the range of variation very tight around 4 (usually 4, occasionally 3 or 5) and very great around 28 (from 14 to 40 taps, for example). Discriminability is a function of the absolute numerical value, as dictated by Weber’s law. Since you were not counting, some other numerical representation must have been guiding your tapping performance – presumably analog magnitudes, as your adherence to Weber’s law, again, would seem to indicate.

Space precludes my reviewing the elegant evidence for analog magnitude representations of number in animals and human infants, but let me give just one example. Fei Xu and Elizabeth Spelke showed infants arrays of dots, one dot array at a time, until the infants got bored with looking at them. All other variables that could have been confounded with number (total array size, total volume of dots, density of dots, and so on) were controlled in these studies, such that the only possible basis for the infants’ discrimination was numeric. Seven-month-old infants were habituated either to arrays of eight or sixteen

dots. After habituation they were presented with new displays containing either the same number of dots to which they had been habituated or the other number. Xu and Spelke found that the infants recovered interest to the new number, and so concluded that they are capable of representing number. Xu and Spelke also found evidence for Weber’s law: infants could discriminate eight from sixteen and sixteen from thirty-two, but not eight from twelve or sixteen from twenty-four.¹

Infants and animals can form analog magnitude representations of fairly large sets, but these representations are only approximate. Analog magnitude representations of number fall short of the representational power of integers; in this system one cannot represent *exactly* fifteen, or fifteen *as opposed to* fourteen. Nonetheless, analog magnitude representations clearly have numerical content: they refer to numerical values, and number-relevant computations are defined over them.

A second system of representations with numerical content works very differently. Infants and nonhuman primates have the capacity to form symbols for individuals and to create mental models of ongoing events in which each individual is represented by a single symbol. Figure 2 shows how, in this system, sets of one, two, or three boxes might be represented. The figure represents three different possibilities for the format and content of the symbols.

1 For an overview of the evidence for analog magnitude representations of number in both nonhuman animals and human adults, see Stanislas Dehaene, *The Number Sense* (Oxford: Oxford University Press, 1997). For evidence in human infants, see Fei Xu and Elizabeth S. Spelke, “Large Number Discrimination in 6-Month-Old Infants,” *Cognition* 74 (2000): B1 – B11.

*Bootstrapping
& the origin
of concepts*

There is one symbol for each box, so number is implicitly represented; the symbols in the model stand in one-one correspondence with the objects in the world.

To give you a feel for the evidence that infants indeed employ such models, distinct from the analog magnitude representations sketched above, consider the following experiment from my laboratory. Ten- to fourteen-month-old infants are shown a box into which they can reach to retrieve objects, but into which they cannot see. If you show infants three objects being placed, one at a time or all at once, into this box, and then allow them to reach in to retrieve them one at a time, they show by their pattern of reaching that they expect to find exactly three objects there. If the infant has a mental representation of a set of two objects (e.g., object, object) that are hidden from view, and the infant sees a new object being added to the set, the infant creates a mental representation of a set of three (object, object, object). Further, computations of one-one correspondence carried out over these models allow the child to establish numerical equivalence and number order (e.g., Have I got all the objects out of the box or are there more?)

So far, this is just another demonstration that infants represent number. However, an exploration of the limits on infants' performance of this task implicates a different system of representation from the analog magnitude system sketched above.

Performance breaks down at four objects. If the infants see four objects being placed into the box and are allowed to retrieve two of them, or even just one of them, they do not reach persistently for the remaining objects. Remember that in the analog magnitude system of representation, success at

Figure 2
Parallel individuation models

Number of boxes	Image	Abstract	Specific
1 box	□	obj	box
2 boxes	□□	obj obj	box box
3 boxes	□□□	obj obj obj	box box box

Note: one symbol for each individual; no symbols for integers.

numerical comparison is a function of the ratios of the numbers being compared, and that the representations can handle sets of objects at least as big as thirty-two. But in this reaching task, infants succeed at ratios of 2:1 and 3:2, but fail at 4:2 and even 4:1; as soon as the set exceeds three, infants cannot hold a model of distinct items in their short-term memory.²

In sum, human infants (and other primates) are endowed with at least two distinct systems of representation with numerical content. Both take sets of individuals as input. One creates a summary analog representation that is a linear function of the number of individuals in the set. This process is noisy, and the noise is itself a linear function of the set size, with the consequence that the representations are merely approximate. For several reasons, this system is too weak to represent the positive integers. For one, there is likely an upper bound to the set sizes that can be represented by analog magnitudes. More importantly,

2 For evidence of the set-size limits on infants' representations of small numbers of objects, see Lisa Feigenson and Susan Carey, "Tracking Individuals Via Object Files: Evidence from Infants' Manual Search," *Developmental Science* 6 (2003): 568–584.

animals and infants cannot discriminate adjacent integer values once the sets contain more than three or four individuals; that is, they cannot represent *exactly* fifteen or twenty-five or forty-nine, or any other large exact integer. Finally, analog magnitude representations obscure one of the foundational relations among successive integers – that each one is exactly one more than the one before. It is this relation, called the successor relation, that underlies how counting algorithms work and provides the mathematical foundation of integer concepts. Since discriminability of analog magnitudes is a function of the ratio between them, the relation between two and three is not experienced as the same as that between twenty-four and twenty-five; indeed, the latter two values cannot really even be discriminated within this system of representation.

The second system – one symbol for each individual – falls even shorter as a representation of integers. There are no symbols for number in this system at all; the symbols in figure 2 each represent an individual object, unlike those in figure 1, which represent an approximate cardinal value. Furthermore, what can be represented in this system is limited in number to sets of one, two, and three.

The count list ('one, two, three...') is a system of representation that has the power to represent the positive integers, so long as it contains a generative system for creating an infinite list. When deployed in counting, it provides a representation of exact integer values based on the successor function. That is, when applied in order, in one-one correspondence with the individuals in a set, the ordinal position of the last number word in the count provides a representation of the cardinal value of the set – of how many individuals it contains. Successive

symbols in the list refer to cardinal values exactly one apart: 5 is 4 plus 1, 6 is 5 plus 1, and so on.

I have argued so far that the count-list representation of number transcends the representational power of both of the representational systems with numerical content that are available to preverbal infants, for these precursors lack the capacity to represent integers. If this is so, it should be difficult for children to come to understand the numerical function of counting.

And so, indeed, it is difficult for children to learn how counting represents number, and details about the partial understanding they achieve along the way constrain our theories of the learning process. In the United States (and every other place where early counting has been studied, including Western Europe, Russia, China, and Japan) children learn to recite the count list as young two-year-olds, and at this age can even engage in the routine of counting – touching objects in a set one by one as they recite the list. But it takes another year and a half before they work out how counting represents number, and in every culture yet studied, children go through similar stages in working out the meanings of the number words in the count list.

First, children learn what 'one' means and take all other words in the list to contrast with 'one,' meaning 'more than one' or 'some.' The behaviors that demonstrate this are quite striking. If you present young two-year-olds with a pile of pennies and ask them to give you one penny, they comply. If you ask for two pennies or three pennies or five pennies, they grab a bunch, always more than one, and hand them over. They do not create a larger set for 'five' than for 'two.' You might suppose that the plural in 'pennies' is doing the work here, but

*Bootstrapping
& the origin
of concepts*

the same phenomenon is observed in China and Japan, even though Chinese and Japanese do not have a singular-plural distinction, and also in the United States when the contrast is between 'one fish,' 'two fish,' and 'five fish.'

Let us call children at this stage of working out the meanings of number words 'one-knowers.' Many other tasks provide additional evidence that one-knowers truly know only the meaning of the word 'one' among all the words in their count list. For example, if you ask a one-knower to tell you what's on a series of cards that contain one, two, or three fish (up to eight fish), they say 'a fish' or 'one fish' for the card with one, and 'two fish' or 'two fishes' or 'two fishies' for all of the other cards. This again indicates a single cut between the meaning of 'one,' which they grasp, and words for the number of individuals in larger sets, which they do not.

After having been one-knowers for about six to nine months, children learn what 'two' means. At this point they can correctly give you two objects if you ask for 'two,' but they still just grab a bunch (always greater than two), if you ask for 'three,' 'four,' 'five,' or 'six.' After some months as two-knowers, they become three-knowers, and some months later induce how counting works.

The performance of children who have worked out how counting works is qualitatively different from that of the one-, two-, and three-knowers in a variety of ways that reflect the conceptual understanding of counting.

To give just one example, in the task in which children are asked to give the experimenter a certain number of items, say four, one-, two-, and three-knowers usually give the wrong number, and the young counters also sometimes make an error. When asked to check by counting

and then to fix the set, counters invariably adjust the set in the right direction, taking an object away if the set is too large or adding one if it is too small. One-, two-, and three-knowers, in contrast, almost always add more to the set – even if they had counted to five or six or seven when they were checking whether it had four – confirming that they really do not understand how counting determines the meaning of number words.

These data suggest that the partial meanings of number words seem to be organized initially by the semantics of quantifiers – the singular-plural distinction and the meanings of words like 'some' and 'a.' If this is right, then we might expect that children learning languages with quantifier systems that mark numerical contrasts differently from English would entertain different hypotheses concerning the partial meanings of number words. They might break into the system differently. And indeed they do.

Consider first classifier languages such as Chinese and Japanese that do not mark the distinction between singular and plural in nouns, verbs, or adjectives. Two independent studies have found that although children in China and Japan learn the count list as young as English-speaking children do, they become one-knowers several months later and are relatively delayed at each stage of the process. Conversely, Russian has a complex plural system in which the morphological markers for sets of two, three, and four differ from those for five through ten. Two independent studies have shown that even Russian one- and two-knowers distinguish between the meanings of the number words 'two,' 'three,' and 'four,' on the one hand, and 'five,' 'six,' 'seven,' and 'eight,' on the other. Unlike the one- and two-knowers

described above, Russian children in the early stages of working out how counting works grab smaller sets when asked to give the experimenter 'two,' 'three,' or 'four' than when asked to give the experimenter 'five' or more, and use larger numbers for larger sets in the what's-on-this-card task.³

These phenomena concerning young children's partial understanding of the meanings of number words support three interrelated conclusions. First, that it is so difficult for children to learn what 'two' means, let alone what 'five' and 'eight' mean, lends support to the claims that preverbal number representations are not representations of integers, at least not in the format of an integer list. Young children – for a full six to nine months before they work out what 'two' means, and a full year and a half before they work out how the count list represents integers – know how to count, know what 'one' means, and know that 'two,' 'three,' 'four,' 'five,' 'six,' 'seven,' and 'eight' represent numbers larger than 'one.' Second, coming to understand how the count list represents numbers reflects a qualitative change in the child's representational capacities; I would argue that it does nothing less than create a representation of the positive integers where none was available before. Finally, a third possible developmental source of natural number representations, in addition to the preverbal systems described above, may be the representations of numbers within natural language quantifier semantics. Of course, natural language quantifiers, other than the number words in the count list itself, do not

3 For a characterization of the early stages of counting in English, see Karen Wynn, "Children's Acquisition of the Number Words and the Counting System," *Cognitive Psychology* 24 (2) (1992): 220 – 257.

have the power to represent natural numbers either.

The problem of the origin of the positive integers arises at two different time scales – historical and ontogenetic. At the dawn of modern anthropology, when colonial officers went out into the French and English colonial worlds, they discovered many systems of explicit number representation that fell short of a full representation of natural number. They described languages that marked number on nouns, adjectives, and verbs, and which had quantifiers like the English 'one,' 'two,' 'many,' 'some,' 'each,' 'every,' and 'more,' but which had no count list. In this vein, the psychologist Peter Gordon has described the language of the Piraha, an isolated Amazonian people. He has shown that in addition to linguistic quantifiers meaning 'one,' 'two,' and 'many,' the Piraha also have access to the nonverbal systems described above (parallel individuation of small sets and analog magnitude representations of large numerosities). Gordon confirms that they have no representations of large exact numerical values.

Anthropologists and archeologists have described intermediate systems of integer representation, short of integer lists, and these intermediate systems provide evidence for a process of cultural construction over generations and centuries of historical time.⁴ Here I concentrate on ontogenetic time. How do three-year-olds do it? How do they create a representational system with more

4 I would recommend the linguist James Hurford's review of this literature to any reader interested in this process. James Hurford, *Language and Number* (Oxford: Basil Blackwell, 1987). For work on the Piraha, see Peter Gordon, "The Role of Language in Numerical Cognition: Evidence from Amazonia" (under review).

*Bootstrapping
& the origin
of concepts*

power than any on which it is built?⁵ In answering this question, I would appeal to bootstrapping processes.

Bootstrapping processes make essential use of the human capacity for creating and using external symbols such as words and icons. Bootstrapping capitalizes on our ability to learn sets of symbols and the relations among them directly, independently of any meaning assigned to them in terms of antecedently interpreted mental representations. These external symbols then serve as placeholders, to be filled in with richer and richer meanings. The processes that fill the placeholders create mappings between previously separate systems of representation, drawing on the human capacity for analogical reasoning and inductive inference. The power of the resulting system of concepts derives from the combination and integration of previously distinct representational systems.

Let's see how this might work in the present case. We must allow the child one more prenumerical capacity – that of representing serial order. This is no problem – young children learn a variety of meaningless ordered lists, such as 'eeny, meeny, miney, mo.'

We seek to explain how the child learns the meanings of the number words – what 'two' means, what 'seven' means – and how the child learns how the list itself represents number – that the cardinal value of a set enumerated by counting is determined by the order on the list, and that successive numbers on the list are related by the arithmetic successor relation.

5 In a forthcoming book I argue that the same bootstrapping process underlying this marvelous feat in childhood also accounts for the development in historical time, but that argument is beyond the scope of this brief paper.

As described above, the child learns the meanings of the first number words as natural language quantifiers. Children learn the meaning of 'one' just as they learn the meaning of the singular determiner 'a' (indeed, in many languages, such as French, they are the same lexical item).

Some months later, 'two' is learned, just as dual markers are in languages that have singular/dual/plural morphology. Languages with dual markers have a different plural affix for sets of two than the affix for sets greater than two. It is as if English nouns were declined 'box' (singular), 'boxesh' (dual), 'boxeesh' (plural). In this system, the suffix 'esh' would apply just when the set referred to contained exactly two items. By hypothesis, children would learn the meaning of the word 'two' just as they would learn the morphological marker 'esh' – if English plural markers worked that way. By extension, some months later, 'three' is learned just as trial markers are in the rare languages that have singular/dual/trial/plural morphology.

In the early stages of being a one-, two-, or three-knower, the child represents other number words as quantifiers, meaning 'many,' where 'many' is more than any known number word. As I will argue below, it is likely that the nonverbal number representations that support the meanings of the known words is the system of parallel individuation (figure 2), with natural language quantification articulation in terms of notions like 'set' and 'individual.'

Meanwhile, the child has learned the count list, which initially has no semantic content other than its order. The child knows one must recite 'one, two, three, four, five,' not 'two, three, one, five, four,' just as one must say 'a, b, c, d, e,' not 'c, a, e, d, b.'

The stage is now set for a series of mappings between representations. Children may here make a wild analogy – that between the order of a particular quantity within an ordered list, and that between this quantity’s order in a series of sets related by additional individuals. These are two quite different bases of ordering – but if the child recognizes this analogy, she is in the position to make the crucial induction: For any word on the list whose quantificational meaning is known, the next word on the list refers to a set with another individual added. Since the quantifier for single individuals is ‘one,’ this is the equivalent to the following induction: If number word X refers to a set with cardinal value n , the next number word in the list refers to a set with cardinal value $n + 1$.

This bootstrapping story provides different answers for how the child learns the meaning of the word ‘two’ than for how she learns the meaning of ‘five.’ According to the proposal, the child ascertains the meaning of ‘two’ from the resources that underlie natural language quantifiers, and from the system of parallel individuation, whereas she comes to know the meaning of ‘five’ through the bootstrapping process – i.e., that ‘five’ means ‘one more than four, which is one more than three...’ – by integrating representations of natural language quantifiers with the external serial ordered count list.

I began by sketching two systems of preverbal representation with numerical content: the analog magnitude system and the system of parallel individuation. You may have noticed that the analog magnitude system played no role in my bootstrapping story. It would be quite possible to imagine a role for this system in a slightly different bootstrapping proposal, and it may be that such a proposal

would be empirically correct, at least for some children. We do know that children come to integrate their integer list with analog magnitudes, such that ‘five’ comes to mean both ‘one more than four, which is one more than three...’ and ‘—————,’ the analog magnitude symbol for the cardinality of a set of five individuals. This integration is undoubtedly very important; bootstrapping provides richer representations precisely through integration of previously distinct systems of representation.

As important as the integration of the integer list representation with analog magnitude representations may be, there is good reason to believe that this integration is not part of the bootstrapping process through which the concept of positive integers is first understood. Research suggests that it is not until *after* children have worked out how the count list represents number – in fact some six months later – that they know which analog magnitudes correspond to which numbers above five in their count list. That finding – along with the fact that the precise meanings of number words are learned in the order ‘one,’ then ‘two,’ then ‘three,’ followed by the induction of how the count list works – leads me to favor the bootstrapping proposal above.

I doubt that anybody would deny that language helps us occupy the distinctive cognitive niche that we human beings enjoy. It is obvious that culturally constructed knowledge is encoded in language and can then be passed on to new generations through verbal communication – you can tell your children something, saving them from having to discover it themselves. Still, this account misses the equally obvious point that children are often unable to understand what we tell them, because they lack the concepts that underlie our words. The problem then becomes accounting for

*Bootstrapping
& the origin
of concepts*

how they acquire the relevant concepts they need to understand what we are telling them.

I have argued that bootstrapping mechanisms provide part of the solution to this problem. In thinking about how bootstrapping might work, we are led to a fuller appreciation of the role of language in supporting the cultural transmission of knowledge. We cannot just teach our children to count and expect that they will then know what 'two' or 'five' means. Learning such words, even without fully understanding them, creates a new structure, a structure that can then be filled in by mapping relations between these novel words and other, familiar concepts. And so eventually our children do know what 'five' means: through the medium of language and the bootstrapping process sketched here they have acquired a new concept.