

The Times

FRIDAY NOVEMBER 05 1999

Regular habits

BY STEVEN PINKER

The fundamental scientific problem raised by language is to explain its vast expressive power. What is the trick behind our ability to fill each others' heads with so many kinds of thoughts? Many language scientists, from Noam Chomsky to his strongest critics, have sought to identify a single mental mechanism that would explain our way with words. I will argue, on the contrary, that language arises not from one trick but from two. They work by different principles, are learned and used in different ways, and may even reside in different parts of the brain. Their border disputes shape and reshape languages over centuries, and make language not only a tool for communication but also a medium for wordplay and an heirloom of endless fascination.

Each "trick" was identified in the nineteenth century by Continental linguists. The first is the principle of the memorized word, which Ferdinand de Saussure called the arbitrary sign. Though a memorized link between a sound and a meaning is rather simple, it can be effective. Since human memory is vast, we can convey a large number of concepts, simply by memorizing sounds that are paired with them. A typical high-school graduate knows around 60,000 words, which works out to a rate of learning a new sound-meaning association approximately every ninety waking minutes starting at the age of one. Also, these entries require little in the way of computation. Given the sound, you can look up the meaning (in comprehension); given a meaning, you look up the sound (in production).

But, of course, we don't just blurt out individual words. We combine them into phrases and sentences, and that brings up the second trick behind language, combinatorial grammar - what Wilhelm von Humboldt called "the infinitive use of finite media". Everyone who speaks a given language has a recipe or algorithm for combining words in such a way that the meaning of the combination can be deduced from the meanings of the words and the way they are arranged. For example, one English rule says that a sentence is composed of a subject - a noun phrase - followed by a predicate - a verb phrase. The verb phrase in turn can consist of a verb, followed by a noun phrase - the object - followed by a sentence, the complement.

The advantage of combinatorial grammar is that by allowing us to combine symbols we can express new combinations of ideas. Journalists say that when a dog bites a man, that isn't news, but when a man bites a dog, that is news. Grammar allows us to convey news, by reshuffling words in particular orders. Moreover, because our knowledge of language is couched in abstract symbols - noun, verb, subject, object - the same rules allow us to talk about a big dog biting a man and a big bang creating a universe.

Another advantage of grammar is that the number of combinations it can generate grows exponentially with the length of the string. If there are, say, 10,000 nouns with which to begin a sentence, and then 4,000 words one can use to continue it, there are $10,000 \times 4,000 = 40$ million two-word beginnings to a sentence, and the number of possible sentences explodes as you continue to add words to the tail of the growing sentence. A final advantage is that human grammars are recursive: a sentence contains a predicate, which can in turn contain a sentence, which can contain a predicate, and so on. That provides an ability to generate structures of arbitrary size, hence an unlimited number of different sentences.

I suggest that the basic design of human language combines the advantages of these two principles. We have a lexicon of words for both common and idiosyncratic entities like ducks and dogs and men, which depends on the psychological mechanism called memory. And we have a set of grammatical rules for novel combinations of entities, for dogs biting men and men biting dogs, which depends on a mental mechanism of symbol combination.

To test this idea, we need a case in which words and rules can express the same idea, but are psychologically, and ultimately neurologically, distinguishable. I believe we do have such a case: regular and irregular inflection.

Verbs in English and in many other languages come in two flavours. Regular verbs such as "walk, walked", "jog, jogged", and "kiss, kissed" are monotonously predictable; all form the past tense by suffixing the stem with "-ed". The regular verbs are open-ended.

English has thousands of existing regular verbs, and new ones are being added all the time. When "to fax" entered common parlance about fifteen years ago, no one had to run to the dictionary to look up its past tense form; everyone knew it was "faxed". Similarly, "flamed", "dissed", "moshed" and "spammed" can all be deduced without having to hear them in the past-tense form.

This productivity is visible even in children. In 1958, Jean Berko Gleason brought some four-year-olds into the lab and said "Here is a man who knows how to wug. He did the

same thing yesterday. He" The children filled in the blank with "wugged", a word they had never heard before, so they could not have memorized it beforehand; they must have generated it on the fly. And in one sense all children are subjects in an experiment like that, because they pass through a stage in which they produce other forms they could not have heard from their parents, forms like "comed", "goed", "bringed", "taked", and "holded".

And that brings us to the second flavour of verbs in English, the irregular verbs, such as "bring, brought", "hit, hit", "go, went", "sleep, slept", "make, made", "ring, rang" and "fly, flew". In contrast to the regulars, the irregulars are unpredictable. The past tense of "sink" is "sank" or "sunk", while the past tense of "cling" is "clung" but not "clang". The past tense of "think" is neither "thank" nor "thunk" but "thought", and the past tense of "blink" is neither "blank" nor "blunk" nor "blought", but is regular, "blinked".

Also unlike the regulars, the irregulars form a closed class. About 180 verbs are irregular in standard English, and there have not been any recent new ones.

All this leads to a simple theory. Irregular verbs are simply pairs of words. Just as we memorize "duck", we memorize "bring" and we memorize "brought", and then we link the two in memory. Regular verbs are generated by a rule, akin to the rule generating sentences out of subjects and predicates. This rule says that a verb in the past tense may be composed of the verb stem plus the suffix "-ed". If a verb does not supply a past-tense form from memory, the regular rule applies by default; that is how children and adults can say things like "wugged" and "faxed" and "spammed", which cannot have been stored in memory beforehand.

Alas, there is a complication for this neat theory: the irregular verbs display patterns. We find families of irregular verbs such as: "keep, kept", "sleep, slept", "feel, felt" and "dream, dreamt"; "wear, wore", "bear, bore", "tear, tore" and "swear, swore"; "string, strung", "swing, swung", "sting, stung" and "fling, flung". This is not what we would expect if the irregular verbs were memorized individually by rote, in which case they could just as easily all be idiosyncratic.

Moreover, these aren't just redundancies in memory; they are occasionally generalized. Occasionally children make errors like "bring, brang", "bite, bote" and "wipe, wope". In the history of language, every once in a while a new irregular verb appears. "Quit" and "knelt" are only about 200 years old (Jane Austen, for example, used "quitted"), and "snuck", which is now standard among Americans and Canadians under fifty, has been in the language only for 100 years. This is especially obvious if you compare non-standard dialects,

which contain forms like "help, holp", "drag, drug" and "bring, brung". Experimental psychologists can even catch people in the act of generalizing an irregular pattern: when Joan Bybee and Carol Moder asked students "What is the past tense of 'to spling'?", many said "splang" or "splung".

How do we account for these patterns and generalizations, which are neither clearly word-like nor clearly rule-like? Two alternatives to the words-and-rules dichotomy have been proposed. Each tries to stretch one of the components to cover the territory ordinarily allotted to the other.

According to the theory of generative phonology from Noam Chomsky and Morris Halle, there are rules all the way down. Just as we have a rule adding "-ed" to form the regular past tense, we have a suite of rules that generate irregular past-tense forms by substituting vowels or consonants. For example, one rule changes "i" to "u" in verbs like "cling, clung".

A problem for this theory is the family resemblance among the verbs undergoing the rule, such as "string, strung", "sting, stung", "fling, flung", "cling, clung". How do you get the rule to apply to them? If you simply link the rule by stipulation to each of the words, you have no explanation for why the words are so similar. Why "string, strung", "sting, stung", and "fling, flung", which share the consonants before and after the "i", and not "fib, fub", "wish, wush," and "trip, trup"? The obvious move at this point is to distil some common pattern out of the set of words that undergo a rule and append it to the rule as a condition. But that does not work either. Say the rule is restricted to apply only to verbs that begin with two consonants and end with "ng". Such a rule would falsely include verbs like "bring" and "spring", which fit the pattern but whose simple past-tense forms are "brought" and "sprang" (not "brung" and "sprung"). At the same time, the rule falsely excludes words like "stick, stuck" and "spin, spun", which miss the condition by a whisker.

The problem is that the words showing an irregular pattern are "family-resemblance" categories in the sense Ludwig Wittgenstein proposed. No set of properties runs through the entire class; rather, patterns of overlapping similarities probabilistically link various subsets.

This led to something completely different: the theory of parallel distributed processing or artificial neural networks from David Rumelhart and James McClelland and their followers. Rather than having words all the way down, this theory has memory associations all the way up. Rumelhart and McClelland devised a neural network model called a pattern-associator memory, which links not an item to an item, but the features of an item - the sounds composing it - to the features of an item. A word is presented to the network by

turning on units corresponding to the word's sounds. The model is trained with examples of a verb and its past-tense form: "sing, sang", "walk, walked", and so on. It records correlations between features of the stem and features of the past-tense form, and that allows it to generalize a pattern to a new verb if it is similar to verbs it has been trained on: once trained on "ring" and "sing" and "spring" and "cling", it automatically generalizes to "spling", because some of the pieces of "spling" occupy the same representational real estate as the pieces of "ring" and "sing". In the same way, it generalizes from "walked" and "talked" to "balked" and "stalked". It generalizes reasonably well, despite not having any distinction between words and rules; a single mechanism handles regular and irregular forms.

Despite the ingenuity of these alternatives, I will present evidence that the traditional words-and-rules model is right after all, though with a twist. Irregulars really are words stored in memory, but memory is not just a list of slots, but is partly associative, a bit like Rumelhart and McClelland's pattern-associator network: features are linked to features, as well as words being linked to words. As a result, similar words reinforce each other and are easier to memorize, and they create a temptation to generalize to new similar words.

But we cannot do without a rule for the regulars. Irregular forms can get away with a pattern-associator memory, because people's use of irregular patterns really is limited by memory: people apply the patterns only to forms that have been memorized or to forms similar to them. But people generalize far more powerfully when it comes to regular forms. The regular inflection can be applied to any word, regardless of its status in memory. As we shall see, that shows that regular inflection is computed by a mental operation that does not need access to the contents of memory, namely a symbol-combination rule that applies to any instance of the symbol "verb". The evidence consists of unrelated circumstances in which memorized forms are not accessed for one reason or another, but people can still apply the regular pattern.

One example is what happens when the memory entry for a word is weak because the word is rare. We know that memory benefits from repetition: the more often you hear something, the better you remember it. If a word is not used very often, its memory trace will be weak. The prediction of the words-and-rules theory is that this should hurt the irregulars, but not the regulars.

The first test of the prediction comes from the statistical structure of the English language. Here is a list of the top ten verbs in English in order of frequency of occurrence in a million-word corpus: "be, have, do, say, make, go, take, come, see, get". Notice that all ten are irregular: be, was; have, had; do, did; and so on. Now, there cannot be a bottom-ten list,

because in the million-word corpus about 800 words are tied for last place, namely, one-in-a-million, the lowest frequency you can measure in a million-word corpus. But the first ten of that list, in alphabetical order, are "abate, abbreviate, abhor, ablate, abridge, abrogate, acclimatize, acculturate, admix" and "absorb". Notice that all ten are regular, as are the vast majority of the uncommon verbs in English.

The explanation is simple. If a word declines in frequency, there may be a generation of children that hasn't properly memorized its irregular past-tense form. Unable to retrieve an irregular form, they default to the regular, and the verb will change to regular for them and for all subsequent generations. If Chaucer were here today, he would say that the past tense of "cleave" is "clove", the past tense of "crow" is "crew", and the past of "chide" is "chid". Old and Middle English had about twice as many irregular verbs as modern English does. Joan Bybee has shown that it is the rarer verbs that have become regular; the ones that are common remain irregular to this day.

You can feel this force of history yourself. Many low-frequency irregulars sound strange to us, and they are slipping out of the language before our ears. Complete this sequence: I stride, I strode, I have --. "Stridden" doesn't sound quite right to most people, presumably because it is not a form you hear every day. Similarly "smite, smote", "slay, slew", "bid, bade", and "forsake, forsook" have a quaint or stilted sound to them, and some are used in regular alternatives such as "slayed". In contrast, low-frequency regulars always sound fine. If I asked you to complete this sequence - I abrogate, I abrogated, I have --, there is nothing particularly strange about "abrogated", presumably because rarity doesn't hurt a word if it is not dependent on memory to begin with.

A nice illustration of this effect comes from idioms and cliches, where people may be familiar with a verb in only one tense. For example, "to forgo" is not terribly common, but does have a certain liveliness in the sarcastic expression "to forgo the pleasure of", as in "You'll excuse me if I forgo the pleasure of watching the video of your wife giving birth". But if you force the cliché into the past tense, you get something strange: "Last night I forewent the pleasure of watching Herb's vacation slides". Likewise, you can say, "I don't know how she can bear that guy", but it's odd to say "I don't know how she bore that guy". You might say "I dig the Doors, man", but "In the sixties, your mother and I dug the Doors" is peculiar.

This contrast is never found with regular past-tense forms, which always sound as good or as bad as their stems. "We can't afford it" comes out as "I don't know how he afforded it"; "She doesn't suffer fools gladly" transforms into "None of them ever suffered fools gladly", both unexceptionable. The rarity of their past-tense forms doesn't hurt them because

their past-tense forms needn't be stored to begin with. They can be computed on the fly by a rule, and inherit whatever sense of naturalness or unnaturalness inheres in the verb stem. The irregulars, in contrast, consist of two entries in memory, which can part company, one familiar, the other unfamiliar.

A second instance in which memory is unhelpful occurs when a new word is difficult to analogize to words in memory because it is unusual in sound. Recall that Bybee and Moder asked volunteers for the past tense of "spling" and other novel words. With "spling", which is similar to existing irregulars such as "cling", "fling", "string" and "sling", most people offered "splang" or "splung". With "krink", which is less similar, only about half produced "krank" or "krunk". And with "vin", which shares only the vowel, hardly anyone suggested "van" or "vun". Sandeep Prasada and I replicated the experiment, but also varied the similarity of novel words to regular forms. "To plip" is similar to many regular verbs such as "clip", "flip", "strip", "nip", "slip", "sip", "tip" and "trip". "To smeej" doesn't rhyme with any English verb root, and "to ploamf" doesn't sound like anything. We presented the words to human subjects and to the trained pattern-associator model. With the irregular-sounding verbs, the model did a reasonably good impersonation of the human being: they both produced many "splangs" and "splungs", fewer "kranks" and "krunks", and very few "vans" and "vuns". But with the regular-sounding words, people and the model diverged: people offered "plipped", "smeejed", and "ploamfed", whereas the model could get only "plipped". "Smeejed" came out as "leefloag", and "frilged" came out as "freezled".

This failure highlights an infirmity of pattern-associator network models: they don't have the computational device called a variable, a symbol that can stand for an entire class regardless of its content, such as the symbol "verb". A pattern-associator can only associate bits and pieces with bits and pieces. The model has nothing to fall back on if a new item doesn't overlap similar, previously trained items, and can only cough up a hairball of the bits and pieces that are closest to the ones that it has been trained on. People, in contrast, reason that a verb is a verb, and, no matter how strange the verb sounds, they can hang an "-ed" on the end of it.

Yet another circumstance showcasing the power of a rule arises when an irregular form is trapped in memory because of the word's grammatical structure. For example, you might say "All my daughter's friends are low-lives"; you wouldn't say "All my daughter's friends are low-lives", even though the ordinary irregular plural of "life" is "lives". Many people refer to more than one Walkman as Walkmans, not Walkmen. Some irregular verbs mysteriously show up in regular garb in certain contexts.

People say "Powell ringed the city with artillery", not "rang", and that a politician "grandstanded", not "grandstood". This immediately shows that sound alone cannot be the input to the inflection system, because a given input, say "life", can come out the other end of the device either as "lifes" or as "lives", depending on something else.

What is that something else? Many language writers have suggested that it is meaning: a semantic stretching of a word dilutes the associations to its irregular past-tense form, causing people to switch to the regular. But this is simply not true. In example after example, a word's meaning can change in large or small ways, and the irregular form sticks to it like glue. For example, if you prefix a word, its irregular form survives: "eat, ate", "overeat, overate" (not "overeated"); similarly we find "overshot", "overdid", "pre-shrank", and so on. If we form a new noun by compounding, any irregular form comes along with it - "workmen" (not "workmans"), stepchildren, milk-teeth, musk-oxen. If you use a noun metaphorically, that too leaves the irregular untouched: "straw men, snowmen, sawteeth, all God's children". And English has hundreds of idioms in which a verb takes on a wildly different meaning, but in all cases it keeps its irregular past-tense form: "Cut a deal" (not "cutted"), "took a leak", "caught a cold", "hit the fan", "blew them off", "put them down", "came off well", "went nuts" and hundreds of others. So it is not enough simply to add a few units for meaning to an associative memory and hope that any stretch of meaning will cut loose an irregular form and thereby explain why people say "low-lifes" and "grandstanded".

A better theory, from linguistic theory, says that headless words become regular. Let me explain. The point of rules of grammar is to assemble words in such a way that one can predict the properties of the new combination from the properties of the parts and the way they are arranged. That is true not just when we string words into sentences, but when we string bits of words into complex words.

Start with the verb "to eat". We combine it with the prefix "over-", to form a new word, "overeat". The scheme for deducing the properties of the new word from its parts is called the "right-hand head rule": take the properties of the rightmost element and copy them up to apply to the whole word. What kind of word is "overeat"? It's a verb, because "eat" is a verb - the verbhood of "eat" gets copied. What does "overeat" mean? It's kind of eating, eating too much. That's because the meaning of "eat" is to eat, and that gets copied to apply to the whole combination. Finally, what's the past tense of "overeat"? It's "overate", because the past tense of "eat" is "ate", and that information gets copied up to the new combination, too.

Another example: Start with the noun "man". Combine it with "work", to produce a new word, "workman". What kind of word is "workman"? It's a noun, because "man" is a noun; the nounhood gets copied. What does "workman" mean? It's a kind of man, a man who does work: the meaning of "man" is passed upstairs. And what is the plural of "workman"? It's "workmen", because the plural of "man" is "men", and that information, too, gets copied.

But there is a family of exceptions: headless words, which don't get their features from the rightmost morpheme. In some compound words, for example, the meaning pertains to something that the rightmost noun has rather than something the rightmost noun is. For example, what is a low-life? A kind of life? No, it is a kind of person, namely, a person who has (or leads) a low life. In forming the word, you have to turn off the right-hand head rule - that is, plug the information pipeline from the root in memory to the whole word - in order to prevent the word from meaning "a kind of life". And if the pipeline is plugged, there is no longer any way for the irregular plural of "life", "lives", to percolate up. That information is sealed in memory, and the regular "add 's'" rule steps in as the default. Other examples include "still-lives" (not "still-lives"), not a kind of life but a kind of painting. We say "sabre-tooths", not "sabre-teeth", because the word refers not to a kind of tooth but to a kind of cat. "Flatfoots" is American slang for policemen; the plural is not "flatfeet" because a "flatfoot" is not a kind of foot. This also solves the mystery of "Walkmans": a Walkman is not a kind of man.

Another example showing off this mental machinery comes from verbs that are based on nouns. We say that the artillery "ringed the city", not "rang", because the verb comes from a noun: "to ring" in this sense means "to form a ring around". To get a noun to turn into a verb, the usual percolation pipeline has to be blocked, because ordinarily the pipeline allows part-of-speech information to be copied from the root to the newly formed word. And that blocked pipeline prevents any irregularity associated with the sound of the verb from applying to the newly formed word. For similar reasons, we say that a politician "grandstanded", not "grandstood", because the verb comes from the noun "play to the grandstand".

To switch to a very different kind of circumstance in which memorized forms are not accessed but regular inflection is applied: childhood. As I mentioned, children frequently make speech errors, things like "We holded the baby rabbits" and "The alligator goed kerplunk". The words-and-rules theory offers a simple explanation: children's memory retrieval is less reliable than adults'. It is based on an uncontroversial fact about the difference between children and adults: children haven't lived as long. Now, among the experiences we accumulate through the years, is hearing the past tense forms of irregular verbs. Since children haven't heard "held" and "came" and "went" very often, they have a weak memory

trace for those forms. Retrieval will be less reliable, and as long as the child has acquired the regular rule, he or she will fill the vacuum by applying the rule, resulting in an error like "comed" or "holded".

Evidence? First, we can show that weak memory is a factor: Gary Marcus and I found that the more often a child's parent uses an irregular in casual speech, the less often the child makes an error in it. Second, the theory explains why children, for many months, produce no errors with these forms - at first they say "held" and "came" and "went", never "holded" and "comed" and "goed". Why does a child wake up one morning and start to say "holded"? Perhaps because that is the point at which the child has just acquired the "-ed" rule. How can we tell? By looking at what children do with regular verbs. Very young children say things like "yesterday we walk", leaving out the past tense altogether. Then they pass from a stage of leaving out the "-ed" more often than supplying it to a stage of supplying it more often than leaving it out. And the transition is exactly at the point at which the first error like "holded" occurs. This is exactly what we would expect if the child has just figured out that the past-tense rule in English is "add '-ed'". Before that, if the child failed to come up with an irregular form, he had no choice but to use it in the infinitive: "Yesterday, he bring . . ."; once he has the rule, he can now fill the gap by over-applying the regular rule, resulting in "bringed".

The final kind of evidence comes from cases in which the memory system is directly compromised by neurological damage or disease. Michael Ullman and I have asked a variety of neurological patients to fill in the blank in items such as "Every day I like to (verb); yesterday, I . . .". We tested patients with "anomia", an impairment in word finding, typically associated with damage to the posterior perisylvian region of the left hemisphere, which leaves the patient in a constant tip-of-the-tongue state. The patient can't get words into speech quickly enough, and resorts to generic fillers like "stuff", "thing", "guy", and "this and that". Patients with anomia can often produce fluent and mostly grammatical speech, suggesting that their mental dictionaries are more impaired than their mental grammars. With such patients, we found that irregular verbs are harder than regulars, which fits the theory that irregulars depend on memory whereas regulars depend on grammar. We also predicted and observed regularization errors like "swimmed", which occur for the same reason that children (who also have weaker memory traces) produce such errors: they cannot retrieve "held" from memory in time.

And the patients are relatively unimpaired in doing a "wug-test" (Today I wug, yesterday I wugged), because that depends on grammar, which is relatively intact.

By symmetrical logic, brain-injured patients with agrammatism - a deficit in stringing words together into grammatical sequences, typically caused by damage to anterior perisylvian regions of the left hemisphere - should show the opposite pattern. They should have more trouble with regulars, which depend on grammatical combination, than with irregulars, which depend on memory. They should produce few errors like "swimmed", and they should have trouble doing the wug-test. And that is exactly what happens.

Despite the identical function of regular and irregular inflection, irregulars are avoided, but the regular suffix is applied freely in a variety of circumstances, from "chided" to "ploamfed" to "lowlifes" to anomia - that have nothing in common except a failure of access to information in memory. These circumstances are heterogeneous and exotic; obviously we don't have separate neural mechanisms that generate regular forms in each of these cases. Rather, they fall out of the simple theory that the rule steps in whenever memory fails, regardless of the reason that memory fails.

These seemingly disconnected phenomena suggest that regular inflection is computed by a mental operation that does not need access to the contents of memory, namely a symbol-combination rule. I believe that this reveals distinct mental mechanisms that implement the two principles of language responsible for its breathtaking expressive power: memory, for the arbitrary sign that underlies words, and symbolic computation, for the infinite use of finite media underlying grammar.

This is an edited extract of the Colin Cherry Memorial Lecture on Communication, given at Imperial College, London, earlier this year. Steven Pinker's "Words and Rules: The ingredients of language" will be published next month.

Copyright 2001 Times Newspapers Ltd. This service is provided on Times Newspapers' standard terms and conditions. To inquire about a licence to reproduce material from The Times, visit the Syndication website.

Talking Point, September 06, 2001