Words and Rules

Steven Pinker

Department of Brain and Cognitive Sciences

Massachusetts Institute of Technology

I am deeply grateful to my collaborators on this project: Alan Prince, Gary Marcus, Michael Ullman, Sandeep Prasada, Harald Clahsen, Richard Wiese, Anne Senghas, Fei Xu, and Suzanne Corkin. Preparation of this paper was supported by NIH Grant HD 18381. Author's address: E10-016, MIT, Cambridge, MA 02139, USA.

Abstract

The vast expressive power of language is made possible by two principles: the arbitrary sound-meaning pairing underlying words, and the discrete combinatorial system underlying grammar. These principles implicate distinct cognitive mechanisms: associative memory and symbol-manipulating rules. The distinction may be seen in the difference between regular inflection (e.g., *walk-walked*), which is productive and open-ended and hence implicates a rule, and irregular inflection (e.g., *come-came*, which is idiosyncratic and closed and hence implicates individually memorized words. Nonetheless, two very different theories have attempted to collapse the distinction; generative phonology invokes minor rules to generate irregular as well as regular forms, and connectionism invokes a pattern associator memory to store and retrieve regular as well as irregular forms. I present evidence from three disciplines that supports the traditional word/rule distinction, though with an enriched conception of lexical memory with some of the properties of a pattern-associator. Rules, nonetheless, are distinct from pattern-association, because a rule concatenates a suffix to a *symbol* for verbs, so it does not require access to memorized verbs or their sound patterns, but applies as the "default," whenever memory access fails. I present a dozen such circumstances, including novel, unusual-sounding, and rootless and headless derived words, in which people inflect the words regularly (explaining quirks like *flied out*, *low-lifes*, and *Walkmans*). A comparison of English to other languages shows that contrary to the connectionist account, default suffixation is not due to numerous regular words reinforcing a pattern in associative memory, but to a memory-independent, symbol-concatenating mental operation.

Words and Rules

Language fascinates people for many reasons, but for me the most striking property is its vast expressive power. People can sit for hours listening to other people make noise as they exhale, because those hisses and squeaks contain *information* about some message the speaker wishes to convey. The set of messages that can be encoded and decoded through language is, moreover, unfathomably vast; it includes everything from theories of the origin of the universe to the lastest twists of a soap opera plot. Accounting for this universal human talent, more impressive than telepathy, is in my mind the primary challenge for the science of language.

What is the trick behind our species' ability to cause each other to think specific thoughts by means of the vocal channel? There is not one trick, but two, and they were identified in the 19th century by continental linguists.

The first principle was articulated by Ferdinand de Saussure (1960), and lies behind the mental dictionary, a finite list of memorized words. A word is an arbitrary symbol, a connection between a signal and an idea shared by all members of a community. The word *duck*, for example, doesn't look like a duck, walk like a duck or quack like a duck, but wee can use it to convey the idea of a duck because we all have, in our developmental history, formed the same connection between the sound and the meaning. Therefore, any of us can convey the idea virtually instantaneously simply by making that noise. The ability depends on speaker and hearer sharing a memory entry for the association, and in caricature that entry might look like this:

(1)

```
  N
  |
 duck
     /d^k/
     (bird that quacks)
```

The entry, symbolized by the symbol at the center (here spelled as English "duck" for convenience), is a three-way association among a sound (/d^k/), a meaning ("bird that quacks"), and a grammatical category ("N" or noun). Though simple, the sheer number of such entries -- on the order of 60,000 to 100,000 for an English-speaking adult (Pinker, 1994) -- allows for many difference concepts to be expressed in an efficient manner.

Of course, we don't just learn individual words. We combine them into strings when we speak, and that leads to the second trick behind language, grammar. The principle behind grammar was articulated by Wilhelm von Humboldt as "the infinite use of finite media." Inside everyone’s head there is a finite algorithm with the ability to generate an infinite number of potential sentences, each corresponding to a distinct thought. The meaning of a sentence is computed from the meanings of the individual words and the way they are arranged. A fragment of the information used by that computation, again in caricature, might look something like this:

(2)

S --> NP VP

VP --> V (NP) (S)

It captures our knowledge that English allows a sentence to be composed of a noun phrase (the subject) and a verb phrase (the predicate), and allows a verb phrase to be composed of a verb, a noun phrase (the object), and a sentence (the complement). That pair of rules is *recursive*: an element is introduced in the right hand side of one rule which also exists as the left hand side of the other rule, creating the possibility of an infinite loop that could generate sentences of any size, such as "I think that she thinks that he said that I wonder whether ...." This system thereby gives a speaker the ability to put an unlimited number of distinct thoughts into words, and a hearer the ability to interpret the string of words to recover the thoughts.

Grammar can express a remarkable range of thoughts because our knowledge of languages resides in an algorithm that combines abstract symbols, such as "Noun" and "Verb," as opposed to concrete concepts such as "man" and "dog" or "eater" and "eaten." This gives us an ability to talk about all kinds of wild and wonderful ideas. We can talk about a dog biting a man, or, as in the journalist’s definition of "news," a man biting a dog. We can talk about aliens landing at Harvard, or the universe beginning with a big bang, or the ancestors of native Americans immigrating to the continent over a land bridge from Asia during an Ice Age, or Michael Jackson marrying Elvis’s daughter. All kinds of unexpected events can be communicated, because our knowledge of language is couched in abstract symbols that can embrace a vast set of concepts and can be combined freely into an even vaster set of propositions. How vast? In principle it is infinite; in practice it can be crudely estimated by assessing the number of word choices possible at each point in a sentence (roughly, 10) and raising it to a power corresponding to the maximum length of a sentence a person is likely to produce and understand, say, 20. The number is 10^{20} or about a hundred million trillion sentences (Pinker, 1994).

Words and rules each have advantages and disadvantages. Compared to the kind of grammatical computation that must be done while generating and interpreting sentences, words are straightforward to acquire, look up, and produce. On the other hand, a word by itself can convey only a finite number of meanings -- the ones that are lexicalized in a language -- and the word must be uniformly memorized by all the members of a community of speakers to be useful. Grammar, in contrast, allows for an unlimited number of *combinations* of concepts to be conveyed, including highly abstract or novel combinations. Because grammar is combinatorial, the number of messages grows exponentially with the length of the sentence, and because language is recursive, with unlimited time and memory resources speakers could, in principle, convey an *infinite* number distinct meanings. On the other hand, by its very nature grammar can produce long and unwieldy strings and requires complex on-line computation, all in service of allowing people to convey extravagantly more messages than they ever would be called upon to do in real life.

Given these considerations, a plausible specification of the basic design of human language might run as follows. Language maximizes the distinct advantages of words and rules by comprising both, each handled by a distinct psychological system. There is a lexicon of words for common or idiosyncratic entities; the psychological mechanism designed to handle it is simply a kind of memory. And there is a separate system of combinatorial grammatical rules for novel combinations of entities; the psychological mechanism designed to handle it is symbolic computation.

How can we test this theory of language design? In particular, how can we distinguish it from an alternative that would say that language consists of a single mechanism that produces outputs of different complexity depending on the complexity of the message that must be conveyed: short, simple outputs for elementary concepts like "dog," and complex, multi-part outputs for combinations of concepts like "dog bites man"? According to the word/rule theory, we ought to find a case in which words and rules express the *same* contents -- but they would still be psychologically, and ultimately neurologically, distinguishable.

I suggest there is such a case: the contrast between regular and irregular inflection. An example of regular inflection can be found in English past tense forms such as *walk-walked, jog-jogged, pat-patted, kiss-kissed*, and so on. Nearly all verbs in English are regular, and the class is completely *predictable*: given a regular verb, its past tense form is completely determinate, the verb stem with the suffix *d* attached.[1] The class of regular verbs is open-ended: there are thousands of existing verbs, and hundreds of new ones being added all the time, such as *faxed*, *snarfed*, *munged*, and *moshed*. Even preschool children, after hearing a novel verb like *rick* in the laboratory, easily create its regular past tense form, such as *ricked* (Berko, 1958). Moreover, children demonstrate their productive use of the rule in another way: starting in their twos, they produce errors such as *breaked* and *comed* in which they overapply the regular suffix to a verb that does not allow it in standard English. Since they could not have heard such forms from their parents, they must have created them on their own. The predictability and open-ended productivity of the regular pattern suggests that regular past tense forms are generated, when needed, by a mental rule, similar in form to other rules of grammar, such as "to form the past tense of a verb, add the suffix *-ed*":

(3)

$$V_{past} \dashrightarrow V_{stem} + d$$

As with other combinatorial products of grammar, regulars would have the advantage of open-endedness, but also the disadvantage of complexity and unwieldiness: some regular forms, such as *edited* and *sixths*, are far less pronounceable than simple English verbs.

In contrast, English contains about 180 "irregular" verbs that form their past tense in idiosyncratic ways, such as *ring-rang*, *sing-sang*, *go-went*, and *think-thought*. In contrast with the regulars, the irregulars are unpredictable. The past tense of *sink* is *sank*, but the past tense of *slink* is not *slank* but *slunk*; 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 regular *blinked*. Also in contrast to the regulars, irregular verbs define a closed class: there are about 180 of them in present-day English, and there have been no recent new ones. And they have a corresponding advantage compared with the regulars: there are no phonologically unwieldy forms such as *edited*; all irregulars are monosyllables (or prefixed monosyllabes such as *become* and *overtake*) that follow that canonical sound pattern for simple English words. The idiosyncrasy and fixed number of irregular verbs suggests that they are memorized as pairs of ordinary lexical items, linked or annotated to capture the grammatical relationship between one word and the other:

---

[1]There are three pronunciations of this morpheme -- the [t], [d], and [ɨd] in *walked*, *jogged*, and *patted*, respectively -- but they represent a predictable phonological alternation that recurs elsewhere in the language. Hence they appear to be the product of a separate process of phonological adjustment applying to a single underlying morpheme, /d/; see Pinker & Prince (1988).

(4)

```
     V          V
     |   ---    |
   bring      brought_past
```

Finally, the memory and rule components appear to interact in a simple way: If a word can provide its own past tense form from memory, the regular rule is blocked; that is why adults, who know *broke*, never say *breaked*. Elsewhere (by default), the rule applies; that is why children can generate *ricked* and adults can generate *moshed*, even if they have never had a prior opportunity to memorize either one.

The existence of regular and irregular verbs would thus seem to be an excellent confirmation of the word/rule theory. They are equated for length and complexity (both being single words), for grammatical properties (both being nonfinite forms, with identical syntactic privileges), and meaning (both expressing the pastness of an event or state). But regular verbs bear the hallmark of rule products, whereas irregular verbs bear the hallmark of memorized words, as if the two subsystems of language occasionally competed over the right to express certain meanings, each able to do the job but in a different way.

The story could end there were it not for a complicating factor. That factor is the existence of *patterns* among the irregular verbs: similarities among clusters of irregular verbs in their stems and in their past tense forms. For example, among the irregular verbs one finds *keep-kept, sleep-slept, feel-felt,* and *dream-dreamt*; *wear-wore, bear-bore, tear-tore,* and *swear-swore*; and *string-strung, swing-swung, sting-stung,* and *fling-flung* (see Bybee & Slobin, 1982; Pinker & Prince, 1988). Moreover, these patterns are not just inert resemblances but are occasionally generalized by live human speakers. Children occasionally produce novel form such as *bring-brang, bite-bote,* and *wipe-wope* (Bybee & Slobin, 1982). The errors are not very common (about 0.2% of the opportunities), but all children make them (Xu & Pinker, 1995). These generalizations occasionally find a toehold in the language and change its composition. The irregular forms *quit* and *knelt* are only a few centuries old, and *snuck* came into English only about a century ago. The effect is particularly noticeable when one compares dialects of English; many American and British dialects contain forms such as *help-holp, drag-drug,* and *climb-clumb*. Finally, the effect can be demonstrated in the laboratory. When college students are given novel verbs such as *spling* and asked to guess their past tense forms, most offer *splang* or *splung* among their answers (Bybee & Moder, 1983).

So the irregular forms are not just a set of arbitrary exceptions, memorized individually by rote, and therefore cannot simply be attributed to a lexicon of stored items, as in the word-rule theory. Two very different theories have arisen to handle this fact.

One is the theory of generative phonology, applied to irregular morphology by Chomsky and Halle (1968) and Halle and Mohanan (1985). In this theory, there are minor rules for the irregular patterns, such as "change *i* to *a*," similar to the suffixing rule for regular verbs. The rule would explain why *ring* and *rang* are so similar -- the process creating the past tense form literally takes the stem as input and modifies the vowel, leaving the remainder intact. It also explains why *ring-rang* displays a pattern similar to *sing-sang* and *sit-sat*: a single set of rules is shared by a larger set of verbs.

The theory does not, however, account well for the similarities among the verbs undergoing a given rule, such as *string, sting, fling, cling, sling*, and so on. On the one hand, if the verbs in this subclass are listed in memory and the rule is stipulated to apply only to the verbs on the list, it is a mysterious coincidence that the verbs on the list are so similar to one another in their onsets (consonant clusters such as *st*, *sl*, *fl,* and so on) and in their codas (the velar nasal consonant *ng*). In principal, the verbs could have shared nothing but the vowel *I* that is replaced by the rule. On the other hand, if the phonological pattern common to the stems in a subclass is distilled out and appended to the rule as a condition, then the wrong verbs will be picked out. Take the putative minor rule replacing *I* by ^, which applies to the *sting* verbs, the most cohesive irregular subclass in English. That rule could be stated as "Change *I* to ^ if and only if the stem has the pattern "Consonant -- Consonant -- *I* -- velar-nasal-Consonant." Such a rule would falsely include *bring-brought* and *spring-sprang*, which do not change their vowels tu ^, and would falsely exclude *stick-stuck* (which does change to ^ even though its final consonant is velar but not nasal) and *spin-spun* (which also changes, even though its final consonant is nasal but not velar). The problem is that the irregular subclasses are *family resemblance categories* in the sense of Ludwig Wittgenstein and Eleanor Rosch, characterized by statistical patterns of shared features rather than by necessary and sufficient characteristics (Bybee & Slobin, 1982).

While generative phonology extends a mechanism suitable to regulars -- a rule -- to capture irregular forms, the theory of Parallel Distributed Processing or Connectionism does the opposite, and extends a mechanism suitable to irregulars -- memory -- to capture various degrees of regularity. The key idea is to make memory more powerful. Rather than linking an *item* to a *item*, one links the *features* of an item to the *features* of another item. Similar items, which share features, are partly superimposed in the memory representation, allowing the common patterns to reinforce each other, and new items that are similar to learned items will activate the shared features and hence inherit the patterns that have been learned previously, allowing for a kind of generalization.

Rumelhart and McClelland (1986) used this principal, which dates back at least to the British associationists' Law of Similarity, to devise a connectionist "pattern associator" memory, with the following major components. The model has a bank of input nodes, each representing a bit of sound of an input stem such as "vowel between two consonants" or "stop consonant at the end of a word." It also has an identical bank of output units representing the past tense form. Every input node is connected to every output node. A verb is presented to the model by first dissolving it into its phonological features and turning on the subset of input nodes that corresponds to the features of the word. These nodes pass activation along the connections to the output nodes, raising their activation to varying degrees. The past tense form is computed as the word that best fits the active output nodes. The activation that is transmitted along the connections depends on the "strengths" of the connections, and these connections are altered gradually during a training phase. Training consists of presenting the network with stems and their correct past tense forms; the connection strengths change gradually to capture the correlations among stem features and past features, averaged over the different stems and pasts in the training set.

For example, the connection between the features in *ing* and the features in *ung* would be strengthened by *cling-clung*, *string-strung*, *stick-stuck*, and so on. Since connections are shared by any verb with given stem features, the trained model can generalize to new verbs according to their similarity to previously trained verbs and according to the strength of the connections from the shared features. The model generalizes because similar verbs are represented in overlapping sets of nodes, so any connection that is trained for one verb is automatically activated by a similar verb. In that way the pattern associator memory, implemented as a computer program, succeeded in learning several hundred regular and irregular verbs, and in generalizing with moderate success to new ones. It did so without any representations specific to words or specific to rules, using instead a single mechanism to handle regular and irregular forms. Subsequent connectionist models (e.g., Plunkett & Marchman, 1991; MacWhinney & Leinbach, 1991; Hoeffner, 1992; Hahn & Nakisa, 1998) differ in their details but share the assumption of a single pattern associator memory for regular and irregular forms, driven by input phonological patterns and generalizing according to phonological similarity.

In this paper I will present evidence that neither of these alternatives to the word-rule theory is called for by the facts of regular and irregular inflection. I will argue for a modified version of the word-rule theory in which irregular forms are still words, stored in memory. Memory, however, is not just a list of unrelated slots, but is partly associative: features are linked to features (as in the connectionist pattern associators), as well as words being linked to words. By this means, irregular verbs are predicted to show the kinds of associative effects that are well-modeled by pattern associators: families of similar irregular verbs are easier to store and recall (because similar verbs repeatedly strengthen a single set of connections for their overlapping material), and people are occasionally prone to generalize irregular patterns to new verbs similar to known ones displaying that pattern (because the new verbs contain features that have been associated with existing irregular families).

On the other hand, I will argue that *regular* verbs are generated by a standard symbol-concatenation rule. I will present evidence that whereas irregular inflection is inherently linked to memorized words or forms similar to them, regular inflection can apply to *any* word, regardless of its memory status. That implies that regular inflection is computed by a memory operation that does not *need* access to the contents of memory, specifically, a symbol-processing operation or rule, which applies to any instance of the symbol "verb." The evidence will consist of a dozen circumstances in which memorized forms are not accessed, for one reason or another, and in which as a consequence irregular inflection is suppressed and regular inflection is applied.

Circumstances in Which Memory is Not Accessed But Regular Inflection Applies

1. Weak Memory Entry (Rare Word). The first circumstance of compromised memory access comes out of the fact that human memory traces generally become stronger with repeated exposure. Thus if a word is rare, its entry in the mental lexicon will be weaker. The prediction of the modified word-rule theory is that irregular inflection will suffer, but regular inflection will not. Several effects of frequency bear this out.

One is the statistical structure of the English language. Here is a list of the ten most frequent verbs in English, in order of decreasing frequency in a million-word corpus (Francis & Kucera, 1982):

```
1. be                 39175/million
2. have               12458
3. do                  4367
4. say                 2765
5. make                2312
6. go                  1844
7. take                1575
8. come                1561
9. see                 1513
10. get                1486
```

Note that all ten are irregular. Compare now the *least* frequent verbs in the corpus, the thousand verbs tied for last with a frequency of one occurrence per million words. The first ten in an alphabetical ordering of them are:

```
3791. abate               1/million
3791. abbreviate          1
3791. abhor               1
3791. ablate              1
3791. abridge             1
3791. abrogate            1
3791. acclimatize         1
3791. acculturate         1
3791. admix               1
3791. adsorb              1
...
```

Note that all ten are regular, as are 98.2% of the rest of the list. (The remainder comprise one irregular root, *smite-smote*, and sixteen low-frequency prefixed versions of high-frequency irregular roots, such as *bethink*, *outdraw*, and *regrind*).

As this comparison shows, there is a massive correlation, in English and most other languages, between token frequency and irregularity. A simple explanation is that irregular forms (but not regular forms) have to be memorized repeatedly, generation after generation, to survive in a language, and that the commonly heard forms are the easiest to memorize. If an irregular verb declines in popularity, at some point a generation of children will fail to hear its past tense form often enough to remember it. Since the regular rule can apply to an item regardless of its frequency, that generation will apply the regular suffix to that verb, and it will be converted into a regular verb for that and all subsequent generations.

Joan Bybee (1985) has gathered evidence for this conjecture. Old English contained about twice as many strong irregular verbs as Modern English, including now-obsolete forms such as *cleave-clove*, *crow-crew, abide-abode, chide-chid, and geld-gelt*. Bybee examined the current frequencies of the surviving descendants of the irregulars in Old English and found that it was the low-frequency verbs that were converted to regular forms over the centuries.

Today one can actually feel the psychological cause of this historical change by considering past tense forms that are low in frequency. Most low-frequency irregulars sound stilted or strange, such as *smite-smote*, *slay-slew*, *bid-bade*, *spell-spelt*, and *tread-trod* (in American English), and one can predict that they will eventually go the way of *chid* and *crew*. In some cases a form is familiar enough to block the regular version, but not quite familiar enough to sound natural, and speakers are left with no clear choice for that slot in the conjugational paradigm. For example, many speakers report that neither of the past participle forms for *stride*, *strided* and *stridden*, sounds quite right. In contrast, low-frequency *regular* past tense forms always sound perfectly natural (or at least no more unnatural than the stems themselves). No one has trouble with the preterite or participle of *abate-abated*, *abrogate-abrogated*, and so on.

An excellent demonstration of this effect comes from cliches, idioms, and collocations, which are often used in a characteristic tense, allowing one to unconfound the familiarity of the *verb* with the unfamiliarity of *a past tense form* of the verb. A verb may be familiar in such a collocation, but only in one tense (say, present or infinitival); by shifting the tense, one can test judgments of a rare tensed form of a common verb. For irregular collocations, the result can often sound strange:

(5)

You will excuse me if I forgo the pleasure of reading your paper before it's published.
*?Last night I forwent the pleasure of grading student papers.

I don't know how she can bear that guy.
*?I don't know how she bore that guy.

I dig The Doors, man!
??In the 60's, your mother and I dug The Doors, son.

That dress really becomes her.
*But her old dress became her even more.

In contrast, expressions containing *regular* verbs that are habitually used in an infinitival form

do not sound any stranger when put in the past tense; they sound exactly as good or bad as their stems. The following examples are common as stems and rare as past tense forms (because they typically occur in negations), but the past tense forms are unexceptionable:
(6)

We can't afford it.
I don't know how he afforded it.

She doesn't suffer fools gladly.
None of them ever suffered fools gladly.

Michael Ullman (1993) has confirmed these claims quantitatively in several questionnaire studies in which subject were asked to rate the naturalness of several hundred regular and irregular forms spanning large ranges in the frequencies of their stem and past tense forms. He found that ratings of regular pasts correlate significantly (.62) with their stems but do not correlate (-.14) with their frequencies (partialing out the ratings of the stem). That is what we would expect if judgments of regulars do not depend on the frequency of any stored past tense forms but are based on forms that may be computed there and then with the rule. The naturalness of the past tense form would instead be inherited from the naturalness of the stem, which is the input to the past tense rule. In contrast, ratings of *irregular* past tense forms correlated less strongly (.32), though still significantly, with their stems, and did correlate significantly (.35) with their frequencies (again, partialing out the ratings of stem). This is what we would expect if the familiarity of an irregular form depended both on how familiar the verb is and how well one remembers its past tense form in particular.

2. Difficult-to-analogize (unusual-sounding) verbs. Many connectionist modelers have pointed out that pattern associator memories can generalize to rare or new verbs based on their similarity to well-learned verbs and on the strength of the connections between the familiar sound patterns and their characteristic past tense forms (see Marcus, Clahsen, Wiese, Brinkmann, and Pinker, 1995, for quotations). People do exactly that for irregular verbs, generalizing an irregular pattern to a ne verb if it is highly similar to an existing family of irregular words. However, people treat regular verbs differently; they apply the regular suffix to *any* new word, regardless of its sound.

Recall that Bybee and Moder (1983) found that subjects generalized irregular patterns to novel stems that were highly similar to existing verbs exemplifying that pattern. The effect depended strongly on the degree of similarity. For forms such as *spling*, which is highly similar in onset, vowel, and coda to existing verbs such as *cling*, *fling*, *sling*, *slink*, *stink*, and *string*, about 80% of the subjects provided irregularly inflected forms such as *splang* or *splung*. For forms such as *krink*, which is similar only in vowel and coda, about 50% of the subjects provided *krank* or *krunk*, and for forms such as *vin*, which shares only a vowel with existing irregulars, only about 20% provided *van* or *vun* -- a classic generalization gradient.

Sandeep Prasada and I replicated these three conditions in Bybee's experiment and added three parallel conditions involving novel verbs likely to receive *regular* inflection (Prasada and Pinker, 1993). Prototypical verbs Verbs like *plip* sound like many existing English regular verbs, such as *clip*, *grip*, *slip*, *snip*, *strip*, and *trip*. Intermediate Verbs like *smaig* do not rhyme with any existing English verb root, and Unusual verbs like *ploamph* are phonologically illicit in English and hence are very dissimilar to existing verbs. We presented the six classes of verbs to subjects and to the trained Rumelhart-McClelland pattern associator model.

For the irregular verbs, the model did a reasonable impersonation of the human beings, showing a generalization gradient in which only the novel verbs that sounded like existing irregular verbs were readily given irregular past tense forms. For the regular verbs, on the other hand, the model and the human diverged. People provided regular forms for unusual-sounding *ploamph* at virtually the same rate with which they provided regular forms for familiar-sounding *plip*. The pattern associator, in contrast, had little trouble with *plipped* but was unable to generate forms such as *ploamphed*. Instead, it produced various blends and random combinations such as *smairf-sprurice*, *trilb-treelilt*, *smeej-leefloag*, and *frilg-freezled*.

The failure is instructive. Pattern associator models, unlike symbol-processing computational architectures, do not have the mechanism of a *variable*, such as "Verb," that can stand for an entire class regardless of its content and that can thereby copy over the phonological material of a stem so that it can be systematically modified to yield a past tense form. Instead the model must be painstakingly trained with items exemplifying every input feature (Marcus, 1999). When a new item exemplifying novel combinations of features is presented, the model cannot automatically copy over that combination; it activates whatever output features are most strongly connected to the features of the new item and generates a chimerical output form from them. (For more technical discussion on this limitation of pattern associator models, see Prasada & Pinker, 1993; Marcus, Pinker, Ullman, Rosen, and Xu, 1992; Marcus et al., 1995; Marcus, 1999.)

3. Information about the irregular form is trapped in memory because of the word's grammatical structure. Linguists, both professional and amateur, have long noted the phenomenon of systematic regularization: some irregular verbs mysteriously call for a *regular* past tense form in certain contexts. Here are three examples:

(7)

All my daughter's friends are low-lifes (*low-lives).

I'm sick of dealing with all the Mickey Mouses in this administration (*Mickey Mice).

Boggs has singled, tripled, and flied out (*flown out) in the game so far.

The phenomenon immediately shows that sound alone cannot be the input to the inflectional system. In the last example, a given sound, such as *fly*, can be mapped onto *flew* and *flown* when referring to birds but *flied* when referring to baseball players. The question is: what is the extra input causing the shift?

Many psychologists, connectionists, and prescriptive grammarians have suggested that the missing input features are *semantic*: when a verb is given a new extended or metaphorical meaning, extra features for that meaning are activated, making the new verb less similar to its predecessor and decreasing activation of the associated irregular form (see Kim, Pinker, Prince, and Prasada, 1991, and Kim, Marcus, Pinker, Hollander, and Coppola, 1994, for more precise versions of this explanation, and detailed discussion of its problems). This explanation, however, is clearly false; most semantic modifications in fact leave an irregular verb's inflectional forms intact. Here are some examples:

(8)

Prefixing: overate/*overeated, overshot/*overshooted, outdid/*outdoed, pre-shrank/*pre-shrinked.

Compounding: workmen/*workmans, superwomen/*superwomans, muskoxen/*muskoxes], stepchildren/*stepchilds, milkteeth/*milktooths.

Metaphor: straw men/*mans, snowmen/*snowmans, God's children/*childs, sawteeth/*sawtooths, six feet/*foots long.

Idiom: cut/*cutted a deal, took/*taked a leak, bought/*buyed the farm, caught/*catched a cold, hit/*hitted the fan, blew/*blowed him off, put/*putted him down, came/*comed off well, went/*goed nuts].

These examples show that merely adding semantic features to a pattern associator, in the hopes that the resulting unfamiliarity of new combinations will inhibit highly trained irregular responses, is unlikely to handle the phenomenon. An explanation that does work comes (with modification) from the linguists Paul Kiparsky (1982)and Edwin Williams (1981): *headless* words become regular.

As with syntax in general, the syntax of words encompasses a scheme by which the properties of a novel combination can be predicted from the properties of its parts and the way they are combined. Consider the verb *overeat*. It is based on the verb root *eat*:

(9)

```
 V
 |
eat
```

The root is then joined with a prefix, yielding the following structure:

(10)

```
                V
               / \
 V      prefix   V
 |         |     |
eat       over  eat
```

The result is a new word that has inherited its properties from the properties of the rightmost morpheme inside it, in this case, *eat*. What syntactic category (part of speech) is *overeat*? It is a verb, just as *eat* is a verb. What does *overeat* mean? It refers to a kind of eating -- eating too much -- just as *eat* refers to eating. And what is its past tense form? *Overate*, not *overeated*, just as the past tense form of *eat* is *ate*, not *eated*. It appears that a new complex word inherits its properties from the properties that are stored in the memory entry of the rightmost morpheme -- the "head" -- including any irregular forms. The pipeline of information from the memory entry of the head at the bottom of the tree to the newly created complex lexical item symbolized by the node at the top of the tree can be schematized as follows:

(11)

```
      X
     / \
        X
       / \
          X
```

This principle, sometimes called the Right-Hand Head Rule, explains all of the preserved irregular forms presented above. For example, the compound noun *workman* is formed by prefixing the noun *man* by the verb *work*:

(12)

```
                N
               / \
  N          V   N
  |          |   |
 man      work man
```

Once again, its properties are inherited from *man*, its rightmost morpheme or head. *Workman* is a noun, just as *man* is a noun; it refers to a kind of man, just as *man* refers to a man, and its plural form is irregular *workmen*, because the plural form of *man* is irregular *men*.

Now here is the explanation for systematic regularization. Some complex words are exceptional in being *headless*. That is, they don’t get their properties -- such as grammatical category or referent -- from their rightmost morpheme. The normal right-hand-head rule must be turned off for the word to be interpreted and used properly. As a result, the mechanism that ordinarily retrieves stored information from the word’s root is inactive, and any irregular form stored with the root is trapped in memory, unable to be passed upward to apply to the whole word. The regular rule, acting as the default, steps in to supply the complex word with a past tense form, undeterred by the fact that the sound of the word ordinarily would call for an irregular form.

Here is how the explanation works for one class of regularizations, compounds whose referent *has* rather than *is* an example of the referent of the rightmost morpheme. For example, a *low-life* is not a kind of life, but a kind of person, namely, a person who has or leads a low life. For it to have that meaning, the right-hand head rule, which would ordinarily make *low-life* mean a kind of life (the semantic information stored in memory with *life*), must be abrogated. With the usual data pipeline to the memory entry for the head disabled, there is no way for the other information stored with *life* to be passed upward either, such as the fact that it has an irregular plural form, *lives*. With the irregular plural unavailable, the regular *-s* rule steps in, and we get *low-lifes*.

(13)

```
                        N
                       / X
                      /   X
   N                 A     N
   |                 |     |
  life              low   life
```

Similar logic explains regularized forms such as *still lifes* (a kind of painting, not a kind of life), saber-tooths (a kind of cat, not a kind of tooth), *flatfoots* (policemen, not feet), *bigmouths* (not a kind of mouth but a person who has a big mouth), and *Walkmans* (not a kind of man, but a "personal stereo"). This effect has been demonstrated in experiments in which four- to eight-year-old children are presented with "has-a" compounds, such as *snaggletooth*, and asked to pluralize them. They provide regular plurals at a significantly higher rate than when they pluralize ordinarily novel compound nouns with irregular roots (Kim et al., 1994).

The same explanation works for a second class of regularizations, eponyms. We hear *Mickey Mouses* because the ordinary noun *mouse* was converted to a distinct syntactic category, that for names, when Walt Disney christened his animated murine hero. (Names are syntactically distinct from common nouns, and hence must bear a different lexical category label, possibly "NP," possibly a category specific to proper names. For present purposes it suffices to symbolize that category simply as "Name". See Marcus et al., 1995, for more detailed discussion.) Then in colloquial speech the name *Mickey Mouse* was converted back into a common noun, *a Mickey Mouse*, referring to a simpleton:

(14)

```
                                             N
                                             |
                        name                name
                         /\                  /\
                        /  \                /  \
   N                 name   N            name   N
   |                   |    |              |    |
 mouse              Mickey  mouse       Mickey mouse
```

The new noun is headless, because the right-hand-head rule had to be turned off twice -- to convert the noun *mouse* into a name, and then to convert the name back into a noun, both violations of the usual upward-copying process. With that process disabled, the irregular plural *mice* remains unexamined in the lexicon, and the regular suffixation rule fills the vacuum and yields *Mickey Mouses*. The same explanation works for most other pluralized irregular-sounding eponyms:[1]

(15)

> The Toronto Maple Leafs/*Leaves (a hockey team named after Canada's national symbol, The Maple Leaf).
> Renault Elfs/*Elves (cars).
> Michael Keaton starred in both Batmans/*Batmen (movie titles).
> We're having Julia Child and her husband over for dinner. You know, the Childs/*Children are really great cooks.

As before, this effect has been shown to work in the language production of children (Kim et al., 1994).

The explanation works in the verb system as well, in the class of regularized past tense and past participle forms of *denominal* verbs: verbs that have been formed out of nouns. In baseball, the verb *to fly* was converted over a century ago into a noun, *a fly*, referring to a high arcing ball. The noun was then converted back into a verb, *to fly out*, meaning "to hit a fly that is caught."

(16)

```
                  V
                  |
         N        N
         |        |
 V       V        V
 |       |        |
fly     fly      fly
```

The verb root *to fly* is thus sealed off from the derived verb *to fly out* at two layers of the structure, the one that converted the verb root to a noun (i.e., failed to copy upwards the information the root's category is "verb"), and the one that converted the noun back into a verb. Among baseball cognoscenti who can sense the *fly ball* in *flying out*, the irregular forms *flew* and *flown* are unable to climb out of the lexical entry for *fly*, and *-ed* applies as the last resort, yielding *flied out*. The same explanation works for other denominals, such as *high-sticked/*high-stuck* (hit with a high stick, in hockey), *grandstanded/*grandstood* (played to the grandstand), and *ringed/*rang the city* (formed a ring around). This regularization process can be documented experimentally in adults' and children's attempts to form past tenses of new verbs (Kim et al., 1991, 1994; see those papers, too, for explanations of apparent counterexamples to this principle).

Similar explanations may be applied to four other kinds of rootless or headless derivation; for a full explanation, see Marcus et al. (1995):

(17)

6. Onomatopoeia: The engine pinged/*pang; My car got dinged/*dang.

7. Quotations: While checking for sexist writing, I found three "man"s/*"men" on page 1.

8. Foreign Borrowing: succumbed/*succame; derided/*derode; chiefs/*chieves; gulfs/*gulves (all borrowed from French or Latin).

9. Artificial concoctions (truncations, acronyms): lip-synched/*lip-sanch (from *synchronize*; Ox's/*Ox-en (hypothetical abbreviation for containers of oxygen).

A *limitation* on regular plurals. Though regular forms can appear in many contexts that are closed to irregulars, there is one circumstance in which the reverse is true: inside compound words. An apartment infested with mice may be called *mice-infested* (irregular plural inside a compound), but an apartment infested with rats is called not **rats-infested* (regular plural inside compound] but *rat-infested* (singular form inside compound), even though by definition one rat does not constitute an infestation. Note that there is no semantic difference between *mice* and *rats* that could account for the grammaticality difference; it is a consequence of sheer irregularity. Similar contrasts include *teethmarks* versus **clawsmarks*, *men-bashing* versus **guys-bashing*, and *purple-people-eater* versus **purple-babies-eater*. In experiments in which subjects must rate the naturalness of novel compounds, Anne Senghas, John Kim and I (Senghas, Kim, & Pinker, 1991) have found that people reliably prefer compounds with irregular plurals, such as *geese-feeder*, over compounds with regular plurals, such as *ducks-feeder*, and that the effect is not a by-product of some confounded semantic, morphological, or phonological difference between regular and irregular plurals.

A simple explanation, based loosely on Kiparsky (1982), might run as follows. Morphological composition of words takes place in several stages. First there is a lexicon of memorized roots, including, according to the word/rule theory, irregular forms. That lexicon supplies the input to rules of regular derivational morphology, which creates complex words (including compounds) out of simple words and morphemes, outputting a stem. Stems are then inputted to a third component, regular inflection, which modifies the word according to its role in the sentence. In simplified form, the architecture of morphology would look like this:

(18)

```
Memorized   -->   Complex   -->  Regular
Roots             Word           Inflection
(including        Formation
irregulars)
```

The word *mice*, stored as a root in the first component, is available as an input to the compounding process in the second component, where it is joined to *infested* to yield *mice-infested*. In contrast, *rats* is not stored as a memorized root in the first component; it is formed from *rat* by an inflectional rule in the third component, too late to be inputted to the compounding rule in the second. Hence we get *rat-infested* but not *rats-infested*.

Peter Gordon (1985) showed that 3-5-year-old children are sensitive to this principle. He asked them questions such as, "Here is a monster who likes to eat *X*. What would you call him?" First he trained them on mass nouns such as *mud*, which don't take a plural, to introduce them to the compound construction, in this case *mud-eater*, without biasing their subsequent answers. Then he tested them by asking what they would call a monster who likes to eat *rats*. The children virtually always said *rat-eater*, not *rats-eater*. In contrast, they frequently called a monster who likes to eat *mice* a *mice-eater* -- and those children who occasionally used the overregularized plural *mouses* in other contexts never used it in a plural such as *mouses-eater*.

In an interesting twist, Gordon checked to see whether children had had an opportunity to learn the distinction by noticing irregular-plural-containing-compounds in their parents' speech, such as *teethmarks*, and simultaneously noticing the *absence* of regular-plural-containing-compounds in their parents' speech, such as *clawsmarks*. He found that *neither* kind of plural is common enough in English for children to have reliably heard them; virtually all commonly used compounds take a *singular* first noun, such as *toothbrush*. Therefore children's sensitivity to the *teethmarks/clawsmarks* distinction is likey to be a product of the innate architecture of their language system, not a product of a tabulation of occurring and non-occurring forms in parental speech.

10. Childhood. Let us return now to the circumstances of impeded memory access and their differential effects on regular and irregular inflection. Recall that children, in their third year, begin to overregularize irregular verbs in errors such as *comed*, *holded*, and *bringed*). Many explanations have been offered in the forty years since these errors have been called to the attention of modern psycholinguists, most portraying the child as a relentless pattern-extractor and regularity-imposer. But these theories founder on the fact that children make these errors in a minority of the opportunities to do so. On average, about 95 percent of children's irregular past tense forms are correct.

Gary Marcus and I (Marcus et al., 1992) proposed the simplest conceivable explanation. The most basic and uncontroversial assumption one can make about children is that they have not lived as long as adults (that is what the word "child" *means*). Among the experiences that one accumulates over the years is hearing the past tense forms of irregular verbs. If children have not heard an irregular form sufficiently often, its memory trace will be weaker than the corresponding trace in adults, and they will retrieve it less reliably and with less confidence (just as adults are less confidence with seldom-heard irregular past tense forms such as *smote*). If the child is at the age at which he or she has acquired the regular past tense rule, then the child will fill the gap by applying the rule to the regular, resulting in an overregularization.

Marcus, Ullman, and I gathered several kinds of evidence for this simple account. One example is a reliable effect of frequency: the more often the parent of a child uses an irregular form, the less often that child overregularizes it. This effect held for all nineteen of the children whose speech we examined, with a mean correlation coefficient of -.33. This is expected if overregularization is an effect of insufficiently reinforced memory entries for irregular forms, and disappears as children hear those forms more and more often and remember them more and more reliably.

A second kind of evidence is a simple explanation of the long-noted phenomenon of "U-shaped development," in which children, for several months, use only correct irregular past tense forms (when they overtly mark the past tense of such verbs at all) before producing their first error. Rumelhart and McClelland had suggested that the onset of overregularization was caused by an increase in the proportion of regular verbs in the child's vocabulary, which would provide the child's pattern-associator with a sudden abundance of inputs strengthening the connections for the regular pattern and temporarily overwhelming the connections implementing each of the irregular patterns. But the evidence is inconsistent with that hypothesis. The proportion of regular tokens remains unchanged in the parental speech directed at children, rather than increasing. The proportion of regular types in the child's vocabulary does increase (as it must, since there is a fixed number of irregulars but an open-ended set of regulars), but not at the right times for the Rumelhard-McClelland hypothesis: there is a negative correlation, not a positive one, over time between children's rate of acquiring new regular verbs and their rate of overregularizing irregular verbs.

Instead, we found that the onset of overregularization errors is best predicted by mastery of the regular rule. Before the first error, children leave regular verbs unmarked most of the time (e.g., *Yesterday we walk*; then there is a a transitional phase in which the child begins to mark these verbs most of the time (e.g., *Yesterday we walked*). It is in this transitional phase that the first overregularization of an *irregular* form is produced. We argue that the tandem development of *walked* and *breaked* comes from a single underlying process, the acquisition of the "add *-ed*" rule, which manifests itself in correct performance where the rule is called for and errors where it is not. Prior to the acquisition of the rule, a child who failed to retrieve *broke* had no choice but to leave it unmarked, as in *Yesterday he break it*; once the child possesses the rule, he or she can mark tense even when memory fails, though the form is incorrect by adult standards.

11 & 12. Disorders of word retrieval in the presence of intact grammar. The final and most direct demonstrations that memory impairment specially affects irregular forms comes from studies of neurological patients whose memory or grammatical systems have been differentially disrupted. Ullman and our collaborators presented a number of patients with a battery of past tense elicitation items of the form, "Everyday I like to *verb*. Yesterday I ________." The verbs were regular, irregular, or novel, and the regular and irregular verbs were equated for frequency and, for a subset of them, pronounceability (e.g., irregular *slept* is similar to regular *slapped*; regular *tried* is similar to irregular *bred*). The prediction is that patients who are more impaired on vocabulary retrieval than on grammatical combination should (1) find irregular forms harder to produce than regular ones, (2) should occasionally produce overregularized forms such as *swimmed* (for the same reason that children do), and (3) should have little trouble producing past tense forms for novel verbs such as *plammed*. Conversely, patients who are more impaired on grammatical combination than on vocabulary retrieval should (1) find regular forms to be harder to produce than irregular ones, (2) should rarely produce overregularized forms, and (3) should have grave difficulty producing past tense forms for novel verbs (Ullman, Corkin, Coppola, Hickok, Growdon, Koroshetz, & Pinker, 1997).

In one case study, we tested a patient with *anomic aphasia*: following damage to left posterior perisylvian regions, he had severe difficulty retrieving words, though his speech was fluent and grammatical. Presumably the subsystem serving vocabulary storage or retrieval was more damaged than the subsystem serving grammatical composition. As predicted, he found irregular verbs harder to inflect than regular verbs (60% vs 89%), made frequent overregularization errors (25% of the opportunities), and was fairly good with novel verbs (84%). In a control case study, we tested a patient with *agrammatic aphasia*: following damage to left anterior perisylvian regions, he had severe difficulty combining words into fluent sentences, but was less impaired at retrieving words. Presumably the subsystem serving grammatical composition was more damaged than the subsystem serving word retrieval. As predicted, he found regular verbs harder to inflect than irregular verbs (20% versus 69%), made no overregularization errors, and was poor at inflecting novel verbs (5%). Similar findings have been obtained by other researchers (Caramazza & Badecker, 1991; Marin, Safran, & Schwartz, 1976).

We found a similar double dissociation when testing patients with neurodegenerative diseases. In Alzheimer's Disease, the most obvious symptom is an impairment in memory, including word retrieval, but in many cases the patients can produce relatively fluent and grammatical speech; this dissociation is thought to be caused by greater degeneration in medial-temporal and temporal-posterior regions of the brain than in the frontal regions. That would predict that these patients would behave like anomic aphasics when producing past tense forms, and indeed they do: the anomic Alzheimer's Disease patients we tested had more trouble producing irregular forms than regular forms (60% versus 89%), made frequent overregularization errors (27%), and were relatively successful in providing past tense forms for novel verbs (84%). The opposite pattern was predicted to occur in Parkinsons' Disease. As a result of degeneration in the basal ganglia, which form an essential part of a circuit involving frontal cortex, Parkinson's Disease patients have many of the symptoms of agrammatism, with less severe impairments in retrieving words from memory. As predicted, the more impaired patients we tested had more trouble with regular verbs than with irregulars (80% versus 88%), had trouble inflecting novel verbs (65%), and never produced overregularization errors.

A Crosslinguistic Validation

All of these comparisons are tainted by a possible confound. An additional property differentiating regular from irregular verbs in English is *type frequency*: regular verbs are the majority in English. Only about 180 verbs in modern English are irregular, alongside several thousand regular verbs. Since pattern associators generalize the majority pattern most strongly, it is conceivable that a pattern associator that was suitably augmented to handle grammatical structure would have the regular pattern strongly reinforced by the many regular verbs in the input, and would come to generalize it most strongly, perhaps in all of the default circumstances I have reviewed.

This is a charitable assumption -- taken literally, theories invoking pattern associators are driven by *tokens* rather than types: the models are said to be learn in response to actual utterances of verbs, in numbers reflecting their frequencies of usage, rather than in response to vocabulary entries, inputted once for each verb regardless of its frequency of usage. Moreover, no pattern associator model yet proposed has plausibly handled the various grammatical circumstances involving headlessness (*flied out*, *Mickey Mouses*, and so on) in which irregular forms systematically regularize. But many connectionist researchers have held out the greater type frequency of regular verbs in English as the main loophole by which future pattern associators might account for the psycholinguistic facts reviewed herein (see Marcus et al., 1995, for quotations). To seal the case for the word/rule theory it would be ideal to find a language in which the regular (default) rule applies to a *minority* of forms in the language. Note that this prediction is an oxymoron according to the traditional, descriptive definition of "regular" as pertaining to the most frequent inflectional form in a language and "irregular" to pertain to the less frequent forms. But I am considering a psycholinguistic definition of "regular" as the default operation produced by a rule of grammatical composition and "irregular" as a form that must be specially stored in memory; the number of words of each kind in the language plays no part in this definition.

One language that displays this profile is German. The past tense is expressed in everyday speech by participles, which come in three forms: strong (involving a vowel change and the suffix -*en*), mixed (involving a vowel change and the suffix -*t*), and weak (involving the suffix (-*t*). The weak forms are analogous (and historically homologous) to English regular verbs. The plural is even more complex, coming in eight forms: four plural suffixes (-*e*, -*er*, -*en*, -*s*, and no suffix), some of which can co-occur with an altered (umlauted) stem vowel. The form that acts most clearly as the default, analogous to English *-s*, is *-s*. This complexity, and various differences in the histories of the two languages, allow us to dissociate grammatical regularity from type frequency (see Marcus et al., 1995, for a far more extensive analysis).

Compare, for example, English *-ed* German *-t*, both "regular" by our definition. Among the thousand most frequent verb types in the languages, approximately 85% of those in English are regular, compared to approximately 45% of those in German. With larger samples of verbs, the gap narrows, but English always shows the higher proportion (see Marcus et al., 1995). But despite regular forms being in a large majority in English and a slight minority in German among the most commonly used verbs, speakers treat them alike. English speakers apply *-ed* to rare verbs, such as *ablated*; German speakers apply *-t* to rare verbs, such as *geloetet* ("welded"). English speakers apply *ed* to unusual-sounding verbs such as *ploamphed*; German speakers apply *-t* to unusual sounding verbs such as *geplaupft*. English speakers apply *-ed* to onomatopoeic forms such as *dinged*; German speakers apply *-t* to onomatopoeic forms such as *gebrummt* ("growled"). English speakers regularize irregular-sounding verbs derived from nouns, such as *flied out*; so do German speakers, in forms such as *gehaust* ("housed"). And just as English-speaking children overregularize irregular verbs in errors such as *singed*, German-speaking children produce corresponding errors, such as *gesingt*.

Plurals provide an even more dramatic comparison. In English, *-s* is applied to more than 99% of all nouns; in German, *-s* is applied to only about 7% of nouns. Despite this enormous difference, the two suffixes behave quite similarly across different circumstances of generalization. In both languages, the *-s* suffix is applied to unusual-sounding nouns (*ploamphs* in English, *Plaupfs* in German) and to names that are homophonous with irregular nouns (*the Julia Childs*, *die Thomas Manns*). The suffix is applied, in both languages, to irregular-sounding eponyms (*Batmans*, *Fausts*) and product names (such as the automobile models *Elfs*, in English, and *Kadetts*, in German). The suffix is also applied in both languages to foreignisms, such as English *chiefs*, borrowed from French, and German *Cafes*, also borrowed from French. Regular suffixes are applied to truncations, such as *synched* in English and *Sozis* and *Nazis* (from *Socialist*) in German). Similarly, in both languages the suffix is applied to quotations: *three "man"s* in English, *drei "Mann"s* in German. Despite the relatively few nouns in German speech taking an *-s*-plural, German-speaking children frequently overregularize the suffix in errors such as *Manns*, analogous to English-speaking children's *mans*. Intriguingly, even the circumstance that tends to *rule out* regular plurals in English, namely the first word in a compound, has a similar effect in German: just as we dislike *rats-eater*, German speakers dislike *Autos-fresser* ("cars-eater"). Many of these effects have been corroborated in experiments with German-speaking adults (Marcus et al, 1995) or children (Clahsen & Rothweiler, 1992; Clahsen, Rothweiler, Woest, & Marcus, 1993; Clahsen, Marcus, & Bartke, 1993), and they have been shown in other languages as well, such as Arabic (McCarthy & Prince, 1990).

These results, combined with a glance at the history of the two languages, provide an interesting insight into *why* regular words form the majority of types in many languages (though not German). In proto-Germanic, the ancestor of English and German, a majority of verbs were strong, the forerunners of today's irregular verbs. There was also a precursor of the weak *-ed/-t* suffix: the "dental suffix," perhaps a reduced form of the verb *do*, which applied to borrowings from other languages and to derived forms, just as it does today. As it happens, the major growth areas in English verb vocabulary over the subsequent centuries was in just these areas. English borrowed rampantly from French (because of the Norman invasion in 1066) and from Latin (because of the influences of the Church and Renaissance scholars); I have estimated that about 60% of English verb roots came from these two languages. English is also notorious for the degree to which nouns can be freely converted to verbs; approximately 20% of our verbs our denominal (Prasada & Pinker, 1993).

Intriguingly, both kinds of these kinds of verbs, once introduced into the language, *had* to be regular on grammatical grounds, because they are rootless and headless. So the standard connectionist account of the correlation between type frequency and regularity may have it backwards. It is not the case that a majority of English verbs are regular, and that causes English-speakers to use the regular suffix as the default. Instead, English-speakers and their linguistic ancestors have used the regular suffix as the default for millennia, and that is *why* the majority of today's English verbs became regular. German, which did not experience a centuries-long domination by a French-speakering elite, and which does not convert nouns to verbs as freely, retained a frequency distribution closer to the ancestral language (see Marcus et al., 1995, for more discussion of the history of the past and plural markers in Germanic languages). Despite these differences in frequency across time and space, the psychology of the speakers remains the same.

Conclusions

We have seen that despite the identical function of regular and irregular inflection, irregular forms are avoided, but the regular suffix is applied freely, in a variety of circumstances (from *gelded* to *ploamphed* to *flied out* to *low-lifes* to anomia) with nothing in common except failure of access to information in memory.

Indeed, these circumstances were deliberately highlighted because they are so *heterogeneous* and *exotic*. Ever since Pinker & Prince's (1988) critique of the original Rumelhart-McClelland pattern associator, many connectionist researchers have responded with models containing ad hoc patches designed specifically to handle one or another of these circumstances (e.g., MacWhinney & Leinbach, 1991; Plunkett & Marchman, 1991; Daugherty & Seidenberg, 1992; Hare & Elman, 1992; Nakisa & Hahn, 1997; for critiques, see Marcus et al., 1992, 1995; Marcus, 1995; Prasada & Pinker, 1993; Kim et al., 1994). But the human brain is not wired with separate innate hardware features dedicated to generating seldom-produced quirks such as *Mickey Mouses* or *three "man"s*, as one finds in many of these models; the phenomena should be consequences of the basic organization of the language system.

The thrust of the argument herein is that within the word/rule theory, the phenomena fall out of the assumption that regular forms are default operations applying whenever memory retrieval fails to provide an inflected form. Regular inflection applies freely in any circumstance in which memory fails because regular inflection is computed by a mental operation that does not *need* access to contents of memory, namely, a symbol-processing rule. Moreover, the comparison with German shows that the applicability of the regular as the default is not caused by the regular pattern being the majority of the child's learning experience. The evidence, then, supports the hypothesis that the design of human language comprises two mental mechanisms: memory, for the arbitrary sign underlying words, and symbolic computation, for the infinite use of finite media underlying grammar.

References

Berko, J. (1958) The child's learning of English morphology. *Word*, *14*, 150-177.

Bybee, J. L. (1985) *Morphology: A study of the relation between meaning and form.* Philadelphia: Benjamins.

Bybee, J. L. & C. L. Moder (1983) Morphological classes as natural categories. *Language, 59*, 251-270.

Bybee, J. L. & D.I. Slobin (1982) Rules and schemas in the development and use of English past tense. *Language*, *58*, 265-289.

Chomsky, N. & Halle, M. (1968) *The sound pattern of English*. New York: Harper and Row.

Clahsen, H. & Rothweiler, M (1992) Inflectional rules in children's grammars: evidence from the development of participles in German. *Morphology Yearbook*, 1-34.

Clahsen, H., Rothweiler, M. Woest, A., & Marcus, G. F. (1993) Regular and irregular inflection in the acquisition of German noun plurals. *Cognition, 45*, 225-255.

Clahsen, H., Marcus, G., and Bartke, S. (1993) Compounding and inflection in German child language. Essex Research Reports in Linguistics, 1:1993, Colchester, England. Also under review, *Language*.

Daugherty, K. & Seidenberg, M. (1992) Rules or connections? The past tense revisited. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. pp. 259-264. Hillsdale, NJ: Erlbaum.

de Saussure, F. (1960) *A course in General Linguistics.* London: Peter Owen.

Francis, N. & Kucera, H. (1982) *Frequency analysis of English usage: Lexicon and grammar*. Boston: Houghton Mifflin.

Gordon, P. (1985) Level-ordering in lexical development. *Cognition, 21*, 73-93.

Hahn, U. & Nakisa, R. (1998) German inflection: Single-route or dual-route? Unpublished manuscript.

Halle, M. & Mohanan, K. P. (1985) Segmental phonology of modern English. *Linguistic Inquiry*, *16*, 57-116.

Hare, M. & Elman, J. (1992) A connectionist account of English inflectional morphology: Evidence from language change. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. pp. 265-270. Hillsdale, NJ: Erlbaum.

Hoeffner, J. (1992) Are rules a thing of the past? The acquisition of verbal morphology by an attractor network. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. pp 861-866. Hillsdale, NJ: Erlbaum.

Kim, J. J., Marcus, G. F., Pinker, S., Hollander, M. & Coppola, M. (1992) Sensitivity of children's inflection to morphological structure. Submitted.

Kim, J. J., Pinker, S., Prince, A. Prasada, S. (1991) Why no mere mortal has ever flown out to center field. *Cognitive Science, 15*, 173-218.

Kiparsky, P. (1982) Lexical phonology and morphology. In I. S. Yang (Ed.) *Linguistics in the morning calm*. Seoul: Hansin, pp. 3-91.

MacWhinney, B. & Leinbach, J. (1991) Implementations are not conceptualizations: Revising the verb learning model. *Cognition, 40*, 121-157.

Marcus, G. F. (1993) Negative evidence in language acquisition. *Cognition, 46*, 53-85.

Marcus, G. F. (1995) The acquisition of inflection in children and multilayered connectionist networks. *Cognition, 56*, 271-279.

Marcus, G. F. (1999) *The algebraic mind.* Cambridge, MA: MIT Press.

Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J. & F. Xu (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development, 57* (4, Serial No. 228).

Marcus, G. F., Brinkmann, U., Clahsen, H., Wiese, R. & Pinker, S. (1995) German inflection: The exception that proves the rule. *Cognitive Psychology, 29*, 189-256.

McCarthy, J. & Prince, A. (1990) Foot and word in prosodic morphology: The Arabic broken plural. *Natural Language and Linguistic Theory*, *8*, 209- 283.

Pinker, S. (1994) *The language instinct.* New York: HarperCollins.

Pinker, S., & Prince, A. (1988) On language and connectionism: Analysis of a Parallel Distributed Processing model of language acquisition. *Cognition, 28,* 73-193.

Plunkett, K., & Marchman, V. (1991) U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition, 38,* 43-102.

Prasada, S. and Pinker, S. . (1993) Generalizations of regular and irregular morphological patterns. *Language and Cognitive Processes*, *8,* 1-56.

Rumelhart, D. & J. McClelland (1986) On learning the past tenses of English verbs. Implicit rules or parallel distributed processing? In J. McClelland, D. Rumelhart and the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition.* MIT Press: Cambridge, MA.

Senghas, A., Kim, J. J., Pinker, S., & Collins, C. (1991) Plurals-inside-compounds: Morphological constraints and their implications for acquisition. Paper presented at the Sixteenth Annual Boston University Conference on Language Development, October 18-20, 1991.

Ullman, M. (1993) The computation and neural localization of inflectional morphology. Doctoral dissertation, Dept. of Brain & Cognitive Sciences, MIT.

Ullman, M., Corkin, S., Coppola, M., Hickok, G., Growdon, J. H., Koroshetz, W. J., & Pinker, S. (1997) A neural dissociation within language: Evidence that the mental dictionary is part of declarative memory, and that grammatical rules are processed by the procedural system. *Journal of Cognitive Neuroscience, 9,* 289-299.

Williams, E. (1981) On the notions "lexically related" and "head of a word". *Linguistic Inquiry, 12*, 245-274.

Xu, F. & Pinker, S. (1992) Weird past tense forms. *Journal of Child Language, 22,* 531-556.

Notes

[1]There are some apparent exceptions, discussed in Marcus et al., 1994, and Kim et al., 1995).