

Rules or Associations in the Acquisition of Morphology? The Frequency by Regularity Interaction in Human and PDP Learning of Morphosyntax

Nick C. Ellis

University of Wales, Bangor, UK

Richard Schmidt

University of Hawai'i at Manoa

When fluent English speakers are asked to produce past tense forms, their latencies are affected by frequency of past tense forms when generating irregular inflections, but not when generating regular ones. This interaction has been used to support hybrid accounts of morphosyntax where regular inflections are computed by an affixation rule in a neurally based symbol manipulating syntactic system, while irregular verbs are retrieved from an associative memory. This article describes adult learning of morphosyntax in a novel language where frequency and regularity are factorially combined. The accuracy and latency data demonstrate frequency effects for both regular and irregular forms early in the acquisition process. However, as learning progresses, the frequency effect on regular items diminishes whereas it remains for irregular items. The regularity by frequency interaction is a natural consequence of the power law of practice and is thus entirely consistent with associative learning processes: Regularity is frequency by another name. Performance of a simple connectionist system, when trained on the same materials, shows a very close correspondence to the human acquisition data.

Requests for reprints should be addressed to N. Ellis, School of Psychology, University of Wales, Bangor, Gwynedd, Wales, LL57 2DG, UK.

We thank Martin Wilson for help with administering the experiment and Gordon Brown and Ernest Lee for useful discussions. This work was assisted by grant R000236099 from the Economic and Social Research Council (UK) to the first author.

A preliminary report of some of this data has previously been reported in Ellis, N.C. & Schmidt, R. (1997). Morphology and longer distance dependencies. *Studies in Second Language Acquisition*, 19, 145–171.

INTRODUCTION

Can human morphological abilities be understood in terms of associative processes, or is it necessary to postulate rule-based symbol processing systems underlying these grammatical skills? This question has generated considerable debate in the literature over the past decade, much of it focusing on the behaviour of “regular” and “irregular” inflectional morphology. There are broadly two contrasting accounts. Dual-processing models (for example Marcus, Brinkmann, Clahsen, Wiese, & Pinker, 1995; Pinker & Prince, 1988; Prasada, Pinker, & Snyder, 1990) take the differences in behaviour of regular and irregular inflections to represent the separate underlying processes by which they are produced: Regular inflections are produced by rules (for example, for the past tense “add *-ed* to a Verb”), while irregular inflections are listed in memory. Associative accounts, whether connectionist (e.g. MacWhinney & Leinback, 1991; Plunkett & Marchmann, 1993; Rumelhart & McClelland, 1986) or schema-network (Bybee, 1995) models, assume that both regular and irregular inflections arise from the same mechanism, a single distributed associative network, with the differences in behaviour being due to statistical distributional factors.

This debate often makes reference to one key behavioural difference between regular and irregular inflections: When people are asked to produce past tense forms, their latencies are affected by frequency of past tense forms when generating irregular inflections, but not when generating regular ones. Prasada et al. (1990) and Seidenberg and Bruck (1990) showed that when fluent native English speakers see verb stems on a screen and are required to produce the past tense form as quickly as possible, they take significantly less time (16–29msec in three experiments) for irregular verbs with high past tense frequencies (like *went*) than for irregular verbs with low past tense frequencies (like *slung*), even when stem frequencies are equated. However, there is no effect on latency of the past tense frequency of regular verbs whose past tense is generated by adding *-ed*.

This lack of frequency effect on regular forms has been taken as evidence that grammar cannot be understood solely in terms of associative mechanisms. Pinker (1991) uses it in support of a hybrid account of morphological inflection: Regular verbs (walk–walked) are computed by a suffixation rule in a neurally based symbol manipulating syntactic system, while irregular verbs (run–ran) are retrieved from an associative memory. Briefly, his explanation is as follows: (1) Irregular inflected forms must be memorised since they do not conform to a rule. A general property of associative memory systems is that there are robust frequency effects: Frequently encountered items are better remembered and faster accessed. Thus, low frequency irregular forms take longer to access than high

frequency ones. (2) Regular inflections are not stored in associative memory, but are generated by a rule-based symbolic system, the time to produce the inflected form simply reflecting (a) the time to access the lemma form, and (b) the time to bind procedurally the regular inflectional affix. Thus, there are no frequency effects on their production latencies. For example, *walk* and *afford* are both quite common in their stem forms, but the past tense form *walked* is much more common than is *afforded*. Nevertheless, a rule-generated account predicts that *afforded* will be produced as quickly as *walked*, since the stem forms, being equally frequent, are equally readily accessed, and it takes a constant amount of time to add an *-ed* ending.

Beck (1995) reports similar regularity by frequency interactions in the latencies of productions of non-native speakers, and thus broadens the application of this account to second language learning. Indeed, the effect is generally cited as key evidence for the existence of symbol-manipulating rules in a specifically linguistic mental module underpinning both first and second language acquisition (Eubank & Gregg, 1995; Pinker, 1991).

It is an elegant and attractive argument, and the latency of production data are indeed consistent with such an account. But there are two problems. The first is that there is a simpler, more parsimonious explanation. In this article we will show that a basic principle of learning, the power law of practice, also generates frequency by regularity interactions. Thus, these behavioural dissociations between “regular” and “irregular” forms are equally consistent with connectionist accounts of morphosyntax. The second problem is that, although these theories are trying to explain both language processing and language acquisition, these particular data come from highly fluent language users. It is difficult to gain an understanding of learning and development from observations of the final state, when we have no record of the content of the learners’ years of exposure to language nor of the developmental course of their proficiencies. If we want to understand learning we must study it directly.

The present report therefore describes adult acquisition of second language morphology using a miniature artificial language (MAL) where frequency and regularity are factorially combined. The accuracy and latency data demonstrate frequency effects for both regular and irregular forms early on in the acquisition process. However, as learning progresses, so the frequency effect for regular items diminishes, although it remains for irregular items. The results thus converge on the end-point described by Prasada et al. (1990). However, they also show how subjects reach this endpoint, with the convergence of performance on high and low frequency regular plurals indexing the rate of acquisition of the regular pattern. We next describe a simple connectionist model which was exposed to the same exemplars in the same order as the human subjects. The results of these simulations closely parallel those of the human learners—there are initially

frequency effects for both the regular and irregular forms, but with increased exposure, so the frequency effect for regular forms is attenuated. Thus, a connectionist system, which has no “rules”, can duplicate this “rule-like” behaviour. Rather, as shown by Plaut, McClelland, Seidenberg, and Patterson (1996) for the case of reading, the frequency by regularity interaction is a natural and necessary result of associative learning processes.

HUMAN LEARNING

If we wish to investigate the effects of input and practice on the acquisition of language structure then we need a proper record of learner input. Yet it is virtually impossible to gather a complete corpus of learners' exposure and production of natural language. How can we ascertain how many types and tokens of regular and irregular inflections have been processed by, for example, learners of English or of German? At best, for natural language, we can only guess by extrapolation of frequency counts from language corpora and unverifiable assumptions about registers. Much of the dispute about the implications of the regularity by frequency effect in morphosyntax centres on such assumptions (Bybee, 1995; Marcus et al. 1995; Plunkett & Marchman, 1991; Prasada & Pinker, 1993, Rumelhart & McClelland, 1986). One way around this is to have people learn a miniature artificial language (MAL) under laboratory conditions.

There is a rich tradition of using MALs to investigate processes of acquisition of native language (Braine, Brody, Brooks, Sudhalter, Ross, Catalano, & Fisch, 1990; Moeser & Bregman, 1972; Morgan, Meier, & Newport, 1987; Morgan & Newport, 1981; Palermo & Howe, 1970; Winter & Reber, 1994) and second- and foreign-languages (MacWhinney, 1983; McLaughlin, 1980; Yang & Givon, 1997). The number of published studies is at least in the hundreds, if not more. This is because MAL experiments have many advantages. They allow: (a) a complete log of exposure to be recorded, (b) accuracy to be monitored at each point, (c) factorial manipulation of the potential independent variables of interest and the teasing apart of naturally confounded effects, and (d) relatively rapid collection of data. But these advantages are bought at the cost of reduced ecological validity: (1) MALs are toy languages when compared to the true complexity of natural language, (2) the period of study falls far short of lifespan practice, (3) laboratory learning exposure conditions are far from naturalistic, and (4) volunteer learners are often atypical in their motivations and demographics. All of these very real problems of laboratory research stem from the sacrifices made necessary by the goals of experimental control and

microanalysis of learning in real time. This is the classic “experimenter’s dilemma”: Naturalistic situations limit experimental control and thus the internal logical validity of research; laboratory research limits ecological validity (Jung, 1971).

In adopting MAL research, we are not denying naturalistic field studies. We might caricature the first as providing valid descriptions of artificial language learning and the latter as providing tentative descriptions of natural language learning. However, the use of a MAL in this study avoids at least three problems that have plagued similar experiments using natural languages (Beck, 1995): (1) Uncertainty whether frequencies derived from corpora accurately represent input to learners, (2) problems attributed to interference from phonological similar items in regular and irregular sets (e.g. *lean*: *lend* or *fly*: *flow*) or derived forms (e.g. *head* as a verb derived from a noun), and (3) evidence from only an advanced stage of learning, forcing reliance on logical argumentation rather than empirical evidence to describe acquisition.

Subjects

Seven monolingual English volunteers for the School of Psychology volunteer panel served as subjects. There were three males and four females. They were aged between 18 and 40 years’ old. They were paid £2.50 per hour for their involvement. They usually worked an hour a day at the experiment.

The Miniature Artificial Language

Moeser and Bregman (1972) criticised the generalisability of MAL experiments which involved subjects listening to strings of words from semantically empty languages, because some syntactic rules that were easily acquired when the MAL referred to a stimulus world were not acquired when it did not. The MAL in the present study therefore incorporated reference. The subjects’ initial task was to learn MAL names for 20 picture stimuli. They were told that they were learning vocabulary in a new language. The pictures, drawn from Snodgrass and Vanderwart (1980), are described in the Appendix, along with the stem form of their MAL names and their corresponding plural forms. Like Braine et al. (1990), we chose MAL names which were suggestive of English cognates in order to make them readily learnable; thus, for example, the MAL words for *umbrella* and *fish* are, respectively, “brol” and “pisc”. To the degree that the task only involves ostensive definition and is not embedded in a larger goal-directed setting, it is acknowledgedly limited as an analogue of natural language vocabulary acquisition. However, it allows clean and precise experimental control whilst providing a reasonable model of ostensive vocabulary

learning that occurs to some considerable degree in L1 and even more so in intentional foreign language learning.

Subjects learned the stem forms before studying the plural forms. In the stem learning phase all items appeared equally often. In the subsequent plural learning phase, in order to maximise the sensitivity of the reaction time (RT) measure, plurality in the MAL was marked by a prefix. Half of the items had a regular plural marker ("bu-"); the remaining 10 items had idiosyncratic affixes, as shown in the Appendix. The use of a prefix inflectional system afforded the additional advantage of minimising transfer effects from the subjects' first language since, although it is found in natural languages like Ndebele, it is quite different from English plural formation. Thus, the MAL was designed with English cognates in order to promote positive transfer of learning of the stem forms, and a very different inflection system in order to exclude any morphological transfer. Frequency was factorially crossed with regularity, with half of each set being presented five times more often. The high and low frequency irregular items were matched for initial phoneme to control voice onset time.

Method

The experiment was controlled by a Macintosh LCIII computer programmed with PsyScope (Cohen, MacWhinney, Flatt, & Provost, 1993). Model pronunciations of the MAL lexis spoken by the first author were recorded using MacRecorder. Subjects' vocal reaction times were measured using a voice key.

Stem Learning. Subjects first learned the stem forms of the MAL. This phase consisted of blocks of 20 trials. In each block every picture appeared once in a randomly chosen order—the subjects' frequency of exposure to all of the stem forms was the same. Each trial consisted of the following sequence: (1) one of the pictures appeared mid-screen for 2sec, (2) if the subject thought they knew the picture name, they spoke it into the microphone as quickly as possible, (3) 2sec after picture onset, the computer spoke the correct name for the picture, (4) the experimenter marked the subject's utterance as correct or not by pressing one of two keys. The dependent variables were thus correctness and RT. These blocks of trials were repeated until the subjects knew the MAL names for the pictures, and could begin uttering them within 2sec of stimulus-onset, to a criterion of 100% correct on two successive blocks. At this point they graduated to the plural learning phase.

Plural Learning. This phase used the same procedures except that each block consisted of 80 trials presented in random order: (1) One presentation

of each of the 20 singular forms as in the preceding phase; (2) five presentations of each of the five high frequency regular (HiFreqReg) plural forms; (3) five presentations of each of the five high frequency irregular (HiFreqIrreg) plural forms; (4) one presentation of each of the five low frequency regular (LoFreqReg) forms; and (5) one presentation of each of the five low frequency irregular (LoFreqIrreg) forms. On the singular trials, just one picture appeared midscreen; on the plural trials, a pair of adjacent identical pictures appeared. This phase continued for several (mean = 4.3, range = 0 to 9) blocks beyond the point at which the learners had achieved 100% accuracy on all plural forms in order to monitor increasing fluency as indexed by RT improvement.

Results

Stem Learning. The stem learning data will only be presented in summary, since the major focus of the experiment lies with the plural forms. Subjects took an average of 9.17 (*SD* 5.93) blocks to achieve the criterion of correctness. Some stem forms were easier to learn than others ($F(19, 2161) = 2.307, P < 0.001$). Particularly easy words included “fant” (92% correct over all trials), “pisc” (85%), and “lant” (78%). Particularly difficult words included “prill” (32%), “charp” (43%), and “breen” (46%). However, for purposes of control, it is important to note that the stem forms of the items that were later allocated in the Plural Learning phase to regular/irregular plural morphology or high/low frequency of exposure did not significantly differ in difficulty of learning at this stage: Regularity [$F(1, 16) = 0.703, ns$]; Frequency [$F(1, 16) = 0.569, ns$]; Regularity \times Frequency ($F(1, 16) = 0.029, ns$).

Plural Learning. Subjects partook of between 13 and 15 blocks of this phase.

The key interest lies with the rate of acquisition of the plural forms. We will first describe analyses of accuracy and then RT. These data are shown in Fig. 1.

ANOVA was used to assess the effects of frequency, regularity, and block. For the main effects of regularity and frequency and their interaction we report additional analyses, which determine the robustness of these effects when separately analysed by subjects and by words. There was a significant effect of frequency on accuracy with the advantage going to the high frequency items [overall analysis, $F(1, 5939) = 431.17, P < 0.001$; by subjects, $F(1, 6) = 56.31, P < 0.005$; by words, $F(1, 16) = 172.00, P < 0.001$]. There was a significant effect of regularity, with the regular plurals being learned better than the irregulars [overall analysis, $F(1, 5939) = 81.52, P < 0.001$; by subjects, $F(1, 6) = 6.64, P < 0.05$; by words, $F(1, 16) = 30.50, P < 0.001$].

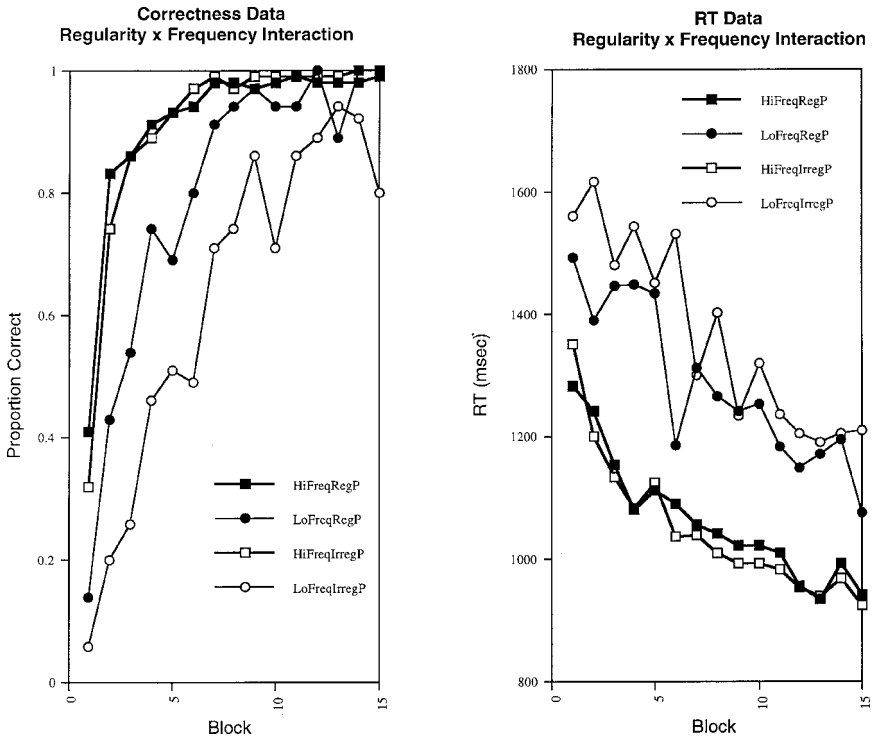


FIG. 1. Acquisition data for human learners of the MAL morphology. The four curves illustrate the interactions of regularity and frequency. The left-hand panel shows accuracy improving with practice. The right-hand panel shows vocal reaction time diminishing with practice. In this graph, as in Figs. 2 and 3, the frequency effect for regular items is assessed by comparing the two solid lines and the frequency effect for irregular items lies in the difference between the two dotted lines.

There was significant improvement over blocks [$F(14, 5939) = 132.00, P < 0.001$]. The interaction of regularity by frequency was significant with the frequency effect being larger for the irregular items [overall analysis, $F(1, 5939) = 73.52, P < 0.001$; by subjects, $F(1, 6) = 12.41, P < 0.02$; by words, $F(1, 16) = 27.73, P < 0.001$]. A significant interaction between regularity by frequency by block [$F(14, 5939) = 2.22, P < 0.005$] shows that the larger frequency effect for irregular items is maximal in the mid-order blocks—it is a lesser effect at early and later stages of learning (Fig. 1).

These patterns are confirmed in the somewhat noisier RT data where the following sources of variation were significant, at least in the overall analysis: (a) frequency [overall analysis, $F(1, 5123) = 650.74, P < 0.001$; by subjects, $F(1, 6) = 63.08, P < 0.001$; by words, $F(1, 16) = 73.96, P < 0.001$]; (b)

regularity [overall analysis, $F(1, 5123) = 10.62, P < 0.001$; by subjects, $F(1, 6) = 3.26, ns$; by words, $F(1, 16) < 1, ns$]; (c) block [$F(14, 5123) = 28.72, P < 0.001$]; (d) regularity by frequency [overall analysis, $F(1, 5123) = 20.92, P < 0.001$; by subjects, $F(1, 6) = 10.15, P < 0.02$; by words, $F(1, 16) = 2.15, ns$]; (e) regularity by frequency by block [$F(14, 5123) = 1.95, P < 0.05$].

It is clear from both panels of Fig. 1 that there was much less regularity effect for high frequency items than for low frequency items and, in counterpart, that the frequency effect was less for regular items. In particular, if the last four blocks of training are taken being typical of more fluent performance, they demonstrate that ceiling effects on the accuracy data allow no frequency effect for the regular items, whereas the effect of frequency is maintained for the irregular ones. The RT curves in the right-hand panel of Fig. 1 are clearly non-linear. In each case a power function better fits the data than does a linear function, the R^2 s for the power function fits being respectively: HiFreqReg 0.94, HiFreqIrreg 0.97, LoFreqReg 0.74, LoFreqIrreg 0.76. Thus the frequency by regularity interaction seems a natural result of asymptotic performance limits: for correctness the 100% accuracy ceiling, for RT the latency "floor" governed by the power law of practice. The curves in Fig. 1 give no hint of a sudden step in performance whereafter all regular items are produced with similar efficiency.

Discussion of Human Data

Like Prasada et al. (1990), these data show a regularity by frequency interaction in the processing of morphology. However, contra Prasada et al., the present data, which concern the *learning* of morphology, demonstrate (a) that there are frequency effects (both on accuracy and RT) for regular items in the early stages of acquisition, (b) the sizes of these effects diminish with learning (converging on a position at fluency as described by Prasada et al.), and (c) the size of the frequency effect on irregular items similarly diminishes with learning, but it does so more slowly.

These effects are readily explained by simple associative theories of learning. It is not necessary to invoke hybrid systems separating rule-governed regular morphosyntax from associatively stored irregulars. If there is one ubiquitous quantitative law of human learning, it is the power law of practice (Anderson, 1982). The critical feature in this relationship is not just that performance, typically time, improves with practice, but that the relationship involves the power law in which the amount of improvement decreases as a function of increasing practice or frequency. Anderson (1982) showed that this function applies to a variety of tasks, including for example, cigar rolling, syllogistic reasoning, book writing, industrial production, reading inverted text, and lexical decision. For the case of language

acquisition, Kirsner (1994) has shown that lexical recognition processes (both for speech perception and reading) and lexical production processes (articulation and writing) are independently governed by the relationship $T=BN^{-\alpha}$ where T is some measure of latency of response and N the number of trials of practice. DeKeyser (1997) shows that automatization of comprehension and production performance involving explicitly learned second-language morphosyntax separately follow independent, skill-specific power functions. Ellis (1996) describes the general implications of the power law for second-language acquisition.

The human acquisition data in Fig. 1 clearly follow the power law of learning. Thus, as performance approaches asymptote, so previously separated functions tend to converge. High frequency items are closer to asymptote. Therefore, whereas performance levels for regular and irregular items are clearly distinguishable at low frequencies, they are much less distinct at high frequencies. This comes as no surprise to us when we consider the ceiling imposed by 100% accuracy. But the power law of practice equally implies an asymptotic ceiling whatever our performance measure.

The power law entails that the contribution of any potential independent variable affecting performance will be more difficult to demonstrate with high-frequency items in practised individuals. This is certainly the case in reading: For example, while spelling and grapheme-phoneme regularity have clear effects on low frequency items, they show little or no effects among high frequency words (Seidenberg, Waters, Barnes, & Tanenhaus, 1984). Our learning data illustrate the same principle operating in the acquisition of morphology. It is not the case that there is no regularity effect on high frequency items (or, concomitantly, no frequency effect on regular items); it is simply that such effects are much smaller closer to asymptote and thus are likely to be swamped by random error. Indeed, high frequency regular inflected forms do exhibit a small (but non-significant) advantage over low frequency forms in naturally occurring errors, and they can be shown to have a larger (significant) advantage in a more controlled experimental task in which subjects produced the past-tense forms of regular English verbs (Stemberger & MacWhinney, 1986).

We have shown that the interaction of frequency and regularity results from developmental trends that are consistent with the ubiquitous descriptive law of associative learning. In the next section we will demonstrate how such data can be generated by a very general mechanism of associative learning. When presented with the same materials at the same relative frequencies of exposure, a standard three-layer feed-forward connectionist model closely simulates our language-learners' acquisition curves.

CONNECTIONIST SIMULATIONS

Connectionist models allow the assessment of just how much of language acquisition can be done by extraction of probabilistic patterns of grammatical and morphological regularities. Since the only relation in connectionist models is strength of association between nodes, they are excellent modelling media in which to investigate the formation of associations (both between surface-form elements and between these and emergent, more abstract, internal representations) as a result of exposure to language. The advantages of connectionist models over traditional symbolic models are that (a) they are neurally inspired, (b) they incorporate distributed representation and control of information, (c) they are data-driven with prototypical representations emerging as a natural outcome of the learning process rather than being prespecified and innately given by the modellers as in more nativist cognitive accounts, (d) they show graceful degradation as do humans with language disorder, and (e) they are in essence models of learning and acquisition rather than static descriptions.

There have been a number of compelling PDP models of the acquisition of morphology. The pioneers were Rumelhart and McClelland (1986), who showed that a simple learning model reproduced, to a remarkable degree, the characteristics of young children learning the morphology of the past tense in English—the model generated the so-called U-shaped learning curve for irregular forms; it exhibited a tendency to overgeneralise, and, in the model as in children, different past-tense forms for the same word could co-exist at the same time. Yet there was no “rule”—“it is possible to imagine that the system simply stores a set of rote-associations between base and past-tense forms with novel responses generated by ‘on-line’ generalisations from the stored exemplars” (Rumelhart & McClelland, 1986, p. 267). This original past-tense model was very influential. It laid the foundations for the connectionist approach to language research which this special issue attests; it generated a large number of criticisms (Lachter & Bever, 1988; Pinker & Prince, 1988), some of which are undeniably valid; and, in turn, it thus spawned a number of revised and improved PDP models of different aspects of the acquisition of the English past tense (e.g. Cottrell & Plunkett, 1994; Daugherty & Seidenberg, 1994; MacWhinney & Leinbach, 1991; Marchman, 1993; Plunkett & Marchman, 1991).

Of these newer models, only that of Daugherty and Seidenberg (1992, 1994) addressed the regularity by frequency interaction. Their model was a three-layer feed-forward network mapping the input of phonological structure of present tense encoded over 120 phonological units representing a CCCVCCC template for English monosyllables onto an output of similarly coded phonological structure of past tense form. Simulation 1,

where the model was trained on all present–past tense pairs with Francis and Kucera frequencies >1 (309 verbs with regular past tenses, 104 verbs with irregular past tenses), failed to generate any frequency by regularity interaction in error score. However, when in simulation 2 the number of irregular verbs in the training set was reduced to just 24, this resulted in there being little effect of frequency on performance with the regular items, whereas performance was better for high frequency irregular verbs than for low frequency ones. This is an important demonstration that the frequency by regularity interaction can be simulated by a connectionist system. However, this model concerned mappings between present- and past-tense forms, not direct access from semantics as in our human data. Furthermore, it is unclear from these simulations how much the results are due to regularity per se, how much to phonological factors (for example, in simulation 1 the error scores for regulars in generalisation tests were inflated by there being a high proportion of phonologically similar irregular past tense false friends in the training corpus, 1994, p. 375), and, given the contrasting results of simulations 1 and 2, how much to the particular choice of training items and the relative proportions of regular and irregular items.

Indeed, much of the debate over the validity of all of these models has concerned (a) the adequacy of the adopted low-level phonological representations, whether these might serve as TRICS (The Representations It Crucially Supposes) which cryptoembody rules within the connectionist network (Lachter & Bever, 1988), (b) over-reliance on phonological cues in models that used sound-to-sound conversions to link base forms with past tense forms (Daugherty & Seidenberg, 1992; MacWhinney, 1994; MacWhinney & Leinbach, 1991), and (c) the appropriateness of the training sets that are used in exposing the models to the evidence of language, and whether they properly reflect the types and tokens, in representative ratios of regular and irregular forms, in a sequence that plausibly mirrors learner language exposure at different stages of development (Daugherty & Seidenberg, 1992; Plunkett & Marchman, 1991). The models are usually concerned with child learner language exposure, yet here the extrapolation is particularly tenuous since adult language frequency norms are typically the only available reference database.

In our simple demonstration, with its intended focus on the frequency by regularity interaction in the acquisition of morphology, we circumvented these problems by the following means:

1. We eliminated TRICS from our input and output representations by entirely ignoring the low level representations and, instead, simply having one input unit for each picture and one output unit for each morpheme. We make no pretence of plausibility of these models for low levels of representation in either input or output processing, but we are presently

neither concerned with low-level feature perception nor the details of motor programming for pronunciation. Each input unit is supposed broadly to correspond to activation of some picture detector or “imagen” (Paivio, 1986), each output unit to some speech output “logogen” (Morton, 1979). We acknowledge that these parts of the model are grossly simplified and we believe that these aspects ultimately involve distributed representations as well. However, there is one advantage to this simplicity—where, as here, each input detector or output logogen is represented by just one unit, with all units having the same form, there is no scope for making some more similar than others, other, that is, than is determined by the frequency of the input–output mappings. This encoding scheme allows the most hygienic investigation of frequency and regularity uncontaminated by other factors.

2. Like Cottrell and Plunkett (1994), we are modelling direct access from semantics rather than generating past tense from stem form phonology. Because there are no phonological representations in our model, there is no chance of the results reflecting any confound with phonology. As usual, costs accompany the benefits: Our simulations can have no bearing on phonological aspects of inflection and thus, while they might generate quantitatively clean data, unlike the elegant error analyses performed by, for example Daugherty and Seidenberg (1994) and MacWhinney and Leinbach (1991), the error responses in the present simulations will be qualitatively uninteresting.

3. We eliminated uncertainty about the detailed content of the complex evidence which human learners are exposed to during their early years of hearing natural language by modelling adult subjects’ learning of the MAL that was reported in the preceding section. Because we determined the exposure sequence of types and tokens of regular and irregular items in this language learning task, we could train the models ensuring the identical history of exposure.

The most common architecture of connectionist model has three layers, the input layer of units, the output layer, and an intervening layer of hidden units (HUs). The presence of HUs enables more difficult input/output mappings to be learned than would be possible if the input units were directly connected to the output units (Broeder & Plunkett, 1994; Rumelhart & McClelland, 1986). The most common learning algorithm is “back propagation” (Rumelhart, Hinton, & Williams, 1986), where, on each learning trial, the network compares its output with the target output, and any difference is propagated back to the hidden unit weights, and in turn to the input weights, in a way that reduces the error. Our simulations adopted this standard architecture. Thus, whatever the pattern of results, they are generated by a very general learning system whose processes were not

tweaked in any way to make it particular as a Language Acquisition Device. So, what are the emergent patterns of language acquisition that result when this general associative learning mechanism is applied to the particular content of picture stimuli with their corresponding singular and plural lexical responses as experienced at the same relative frequencies of exposure as our human learners?

The Models

Architecture. Every model had 22 input (I-) units. Each of I-units 1–20 represented one of the pictures used in the training set of the Appendix. I-unit 21 represented another picture (the generalisation test item, TesterP) which was only ever presented for training to the model in the singular—later it was presented as a plural test item to see which plural affix the model would choose for this generalisation item (akin to asking you what is the plural of a novel word like “wug”?). I-unit 22 coded plurality, that is, whether a singular stimulus item or a pair were presented. Every model had 32 output (O-) units. O-units 1–20 represented the stem forms of the lexis shown in the Appendix. O-unit 21 represented the stem form corresponding to I-unit 21. O-units 22–31 represented each of the other 10 unique plural affixes for irregular items. O-unit 32 represented the regular plural affix. This numbering of I- and O-units is, of course, arbitrary and was randomised across models—what mattered and remained constant was that the same O-unit was always reinforced whenever a particular I-unit was activated.

We investigated four different classes of model, which differed in their computational capacity or resources. The larger the number of HUs in a model, the larger the number of connections in the network, and the greater its capacity to learn new associations and abstractions. Thus, we compared models with 3, 5, 8, and 15 HUs.

Stem Training. At the outset, the connection weights of the models were randomised. Then, just like our human learners, the models were first trained on the singular forms. Each epoch of training consisted of 21 trials. Each trial consisted of presentation of a unique input pattern, one for each of the input pictures. Thus, just one of I-units 1–21 would be “on” on any trial. Throughout the singular training phase, I-unit 22 (representing single/plural stimuli) was set to “off”. For each input pattern, the model responded with a pattern of output over its 32 O-units. Initially this was the random result of the random connection weights. But the model was also presented with the correct pattern of output for that corresponding input pattern (e.g. if I-unit 1

was on and all others off, O-unit 1 should have had value 1.0 and all others zero). On each trial, the back-propagation algorithm calculated the difference between the level of activity that was produced on each O-unit and the “correct” level of activity, and a small adjustment was made to the connection strength to that unit in such a way that when the same process occurred again, a closer approximation to the correct pattern of output activation would be achieved. The models were trained for 500 epochs of singular experience. For each size of model we ran five examples starting with different arbitrary unit allocation and different initial random connection strengths. The data we produce for each model is the average performance of these five examples.

Plural Training. The model weights that resulted from this singular training then served as the starting point for another 700 epochs of training on plurals. The trials constituting each epoch were very similar in nature to those used with the human learners: Each epoch consisted of 81 trials presented in random order: (a) One presentation of each of the 21 singular forms as in the preceding phase, (b) five presentations of each of the five high frequency regular (HiFreqReg) plural forms, (c) five presentations of each of the five high frequency irregular (HiFreqIrreg) plural forms, (d) one presentation of each of the five low frequency regular (LoFreqReg) forms, and (e) one presentation of each of the five low frequency irregular (LoFreqIrreg) forms. For training trials of type (a) just one of I-units 1–21 was activated, I-unit 22 was off, and just the corresponding one of O-units 1–21 was reinforced. For the other training types (b–e), one of I-units 1–20 was activated, I-unit 22 was on, and one of O-units 1–20 (the corresponding stem form) along with one of O-units 22–32 (the corresponding plural affix) were reinforced. The learning algorithm operated as it did in the stem training phase. At regular intervals, we tested the state of learning of the model by presenting it, without feedback, with test input patterns that represented the plural cases of all 21 pictures. At these tests, for each stimulus we measured the pattern of activation (between 0 [no activation], and 1 [full on]) across O-units 22–32 and compared it against the target plural activation for that input pattern.

Results

Regularity by Frequency. Figure 2 shows the Root Mean Square (RMS) error calculated across the plural affix O-units (22–32) averaged over the five items in each of the following classes: HiFreqReg, HiFreqIrreg, LoFreqReg, LoFreqIrreg, at each point in testing of the model. These graphs illustrate that learning in all of the models showed clear effects of frequency (high




FIG. 2. Acquisition data for four connectionist models with increasing computational power trained on the MAL morphology. There are clear regularity by frequency interactions in all models.

frequency items were learned faster than low frequency ones), regularity (regular items were learned faster than irregular ones), and a frequency by regularity interaction whereby there was much less regularity effect for high frequency items than for low frequency items and, equally, that the frequency effect was less for regular items than for irregular ones.

ANOVAs on these RMS data for each size of model demonstrated that there was high consistency of response across items and example simulations. For example, when the 8HU model was analysed as a repeated measures ANOVA across 15 roughly equally spaced blocks of training (to parallel the human data analysis), the following significant effects were observed: (a) Frequency [by simulations, $F(1, 16) = 20.80$, $P < 0.0005$; by words, $F(1, 16) = 56.65$, $P < 0.0001$], (b) regularity [by simulations, $F(1, 16) = 9.07$, $P < 0.01$; by words, $F(1, 16) = 39.57$, $P < 0.0001$], (c) regularity by frequency [by simulations, $F(1, 16) = 4.85$, $P < 0.05$; by words, $F(1, 16) = 15.61$, $P < 0.005$], (d) block [by simulations, $F(14, 224) = 68.03$, $P < 0.0001$; by words, $F(14, 224) = 149.14$, $P < 0.0001$], (e) block by regularity [by simulations, $F(14, 224) = 36.75$, $P < 0.0001$; by words, $F(14, 224) = 29.29$, $P < 0.0001$], (f) block by frequency [by simulations, $F(14, 224) = 18.93$, $P < 0.0001$; by words, $F(14, 224) = 11.84$, $P < 0.0001$], and (g) block by regularity by frequency [by simulations, $F(14, 224) = 16.11$, $P < 0.0001$; by words, $F(14, 224) = 13.06$, $P < 0.0001$].

Comparison of this pattern of ANOVA effects with that reported earlier for the human data shows important similarities: in both cases there are significant main effects of frequency, regularity and blocks, and there are significant interactions involving regularity by frequency and regularity by frequency by block. Thus, the connectionist models demonstrate effects which broadly parallel those found in humans.

Comparison with Human Data. More detailed comparison is also possible. Although RMS error is the usual measure of model performance because it assesses how well the network learns to inhibit non-relevant units as well as to excite relevant ones, we also extracted simple accuracy data for the 8HU model. This accuracy score is the amount of activation (between 0 and 1) on the single O-unit which corresponds to the appropriate target affix for that input pattern. Figure 3 shows the performance of the 8HU model using this metric. It is clear that accuracy scores generate a graph which is effectively a reflection in a horizontal plane of the RMS data shown in the third panel of Fig. 2. In fact, in the current simulations correct activation is almost perfectly correlated with MSE (for example, $r = -0.988$ for the 8HU model). However, the activation metric has the advantage of more ready interpretation and direct comparison with the human data.

When the 8HU model and the human data are aligned as in Fig. 3 these correspondences become clear. Pairwise comparison of individual points

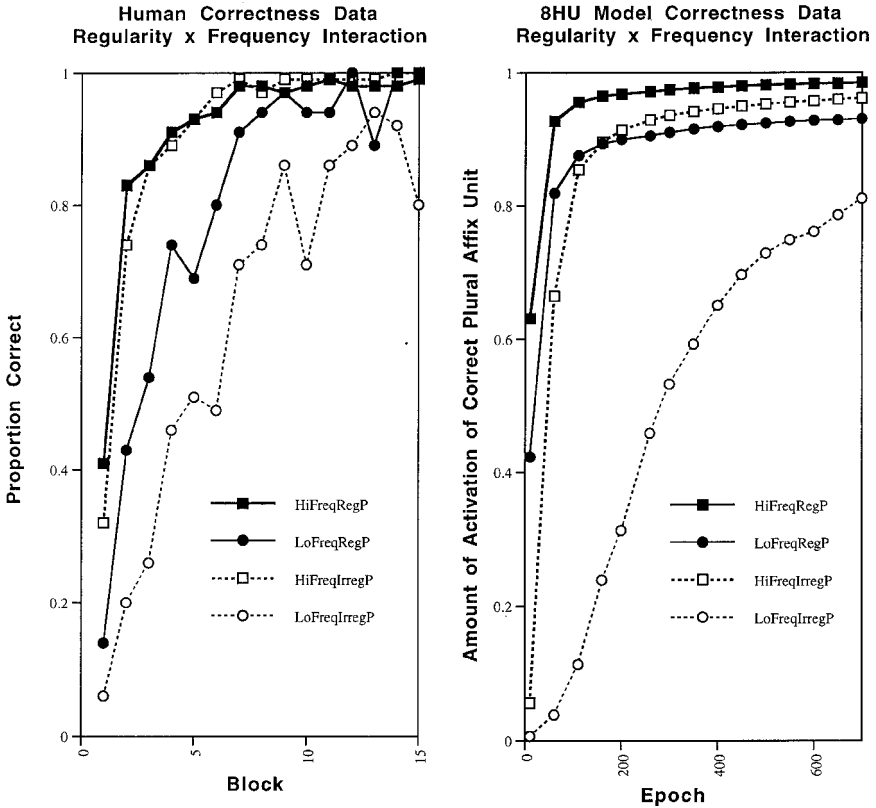


FIG. 3. A comparison of human accuracy performance and that of the eight hidden unit connectionist simulation.

across these two graphs by correlation shows that the simulation predicts a large proportion of the variance in the human data ($R^2 = 0.78$). There are some differences in detail—as is clarified in Fig. 4 where performance is averaged over blocks, the model performs somewhat better on the regular items and worse on the irregular items, particularly the low frequency irregular items, than do the humans. ANOVA (three factor [human/model, regularity, and frequency] with 15 blocks as repeated measures, by words analysis) comparing the human and 8HU model data confirms these interactions: (a) human/model $F(1, 32) = 1.36$, ns, (b) human/model by frequency $F(1, 32) = 0.47$, ns, (c) human/model by regularity $F(1, 32) = 30.28$, $P < 0.0001$, (d) human/model by regularity by frequency $F(1, 32) = 5.01$, $P < 0.05$.




FIG. 4. The regularity by frequency interaction averaged over blocks in humans and the eight hidden unit model. Error bars reflect 95% confidence intervals.

Generalisation. So far we have described performance with trained items. However, we also tested model output when the stimulus was the pattern for generalisation item (TesterP) along with activation of the plural marking I-unit 22, a state of input on which the models had never been trained. Table 1 shows performance of the different models at the end of training. It is clear that the larger models have abstracted the regular plural pattern and tend to apply it by default to the generalisation test item: for the 15HU model, (a) average activation on the regular plural O-unit is 0.60, (b) mean RMS error comparing observed activation across O-units 22–32 and the target regular plural pattern (1000000000) is just 0.45, and (c) four out of the five exemplar runs of this size of model chose the regular plural pattern as being the closest to observed output as measured by minimum

TABLE 1

Performance on the Target Regular Plural Pattern for the Four Sizes of Model When Presented with the Generalisation Wug-test Item, TesterP, at End of Training

<i>Measure</i>	<i>Model Size</i>			
	<i>3HU^a</i>	<i>5HU</i>	<i>8HU</i>	<i>15HU</i>
RMS error ^b				
<i>M</i>	0.81	0.79	0.53	0.45
<i>SD</i>	0.43	0.50	0.45	0.32
Activation weight ^c				
<i>M</i>	0.20	0.28	0.57	0.60
<i>SD</i>	0.44	0.44	0.52	0.35
N. hits (/5) ^d	1	2	3	4

There were five exemplars of each size of model. ^aHU = hidden units. ^bRMS error calculated against the target activation pattern across O-units 22–32 for the regular plural affix. ^cActivation weight on the regular plural affix O-Unit. ^dNumber of exemplar models (/5) which chose the regular plural affix pattern for TesterP as indexed by output weights on O-units 22–32 being closest to the regular plural affix target pattern activation using a squared Euclidean distance metric.

squared Euclidean distance. Thus, when the larger models are presented with a plural stimulus which they have only ever previously experienced as a single form, there is a tendency for them to generalise and apply the regular plural morpheme (bu-) in the same way that humans might generalise that the plural of “wug” is “wugs”.

Effects of Different Sizes of Model. Figure 2 also illustrates the effects of manipulating computational capacity of model: (1) Models with lower computational power (= a smaller number of HUs) learn the high frequency items quite well—almost as well as the largest model. (2) The most striking effect of varying the computational power of the models lies in their abilities to learn low frequency irregular items—this is by far the most sensitive index of morphological learning ability. The 3HU model hardly manages to learn these forms at all. The 15HU model eventually learns them rather well. (3) There is essentially no frequency effect for regular items in the higher computational power models, but none the less the frequency effect for irregular items remains strong. (4) The smaller models continue to show a frequency effect for regular items at the end of training. Table 1 provides one additional effect of model size: (5) The greater the computational power of the models, the more they operate in “rule-like” way by abstracting a “regular” plural form, which is applied by default to novel items. In sum, while lower computational power models are reasonably good on high frequency regular items, they show frequency effects for irregular and

regular items, are relatively poor on “wug tests”, and have particular difficulty on low frequency irregular items.

Discussion of Simulations

We believe that, at least for the issue of regularity and frequency effects in morphosyntax, this is to date the most complete quantitative analysis of the adequacy of fit of simulation to human data. We are not simply making predictions about how an underspecified model might behave (the Daugherty & Seidenberg, 1994 criticisms of the Pinker & Prince, 1988 and Pinker, 1991 theories). We are not simply demonstrating that simulation and human data alike exhibit first order interactions of frequency and regularity (Daugherty & Seidenberg, 1994). Instead we are showing the parallel patterns of significance of main effects, first, and second order interactions in ANOVAs of simulation and human data, and we are showing that the simulations explain close to 80% of the relevant human data. When we go as far as actually comparing human and model performance in a multifactorial ANOVA we find some differences of detail in the size of interactions that are qualified by the human/model factor. But these differences of detail do not detract from the general success of the models in simulating the human pattern of development of the frequency by regularity interaction: In humans and models alike, high frequency items were learned significantly faster than low frequency ones; regular items were learned significantly faster than irregular ones; there was a significant frequency by regularity interaction where the frequency effect was less for regular items than for irregular ones; and this is qualified as the higher level interaction with block whereby there is a developmental trend—the frequency effect for regular items attenuates faster than that for irregular items.

We have demonstrated that the models can generalise and produce the default plural affix for a novel stimulus. Similar “wug test” performance by a human learner would be taken as an operationalisation that they had acquired the “regular” morphological systematicity.

Finally, we have shown how varying the computational capacity of the models affects both the rate of acquisition of default case, the presence or absence of frequency effects for regular items, and ability to acquire irregular items. This is compatible with existing data for children with specific language impairment (SLI). Oetting and Rice (1993) compared five-year-old SLI children with age-matched controls on their ability to form plurals. The SLI children were significantly worse at generating regular plurals for nonce (=wug) items; they were worse at generating regular plurals; and they showed an effect of frequency on the regular items which the control children, because of ceiling effects, did not. Unfortunately, Oetting and Rice (1993) do not provide clear data on the children’s ability to

form irregular plurals. However, their pattern of differences between SLI and control children's performance on regular items is sufficiently close to that between the present low-capacity and high-capacity simulations to suggest that morphosyntactic impairments in individuals with SLI might be explained by reduced language processing capacity in a general associative memory network rather than by a hybrid account. The SLI children's showing frequency effects for regular items is particularly compelling in this respect. However, further assessment of regularity by frequency effects and default abstraction in individuals with SLI and with Williams syndrome (whose ability on regular forms is said to outstrip their performance on irregulars—Bellugi, Bihle, Jernigan, Trauner, & Dougherty, 1990) is necessary to test these parallels further (see Marchman, 1993 for other simulations of different types of language dysfunction).

GENERAL DISCUSSION

Fluent language users have processed many millions of utterances involving tens of thousands of types presented as innumerable tokens. It should come as no surprise either that they demonstrate such effortless and complex skill as a result of this mass of practice, or that researchers, lacking any true record of the learners' experience, are awed and confused by these sophisticated grammatical abilities. While we have no wish to deny any of the complexity of the final fluent state, we suspect that much of the mystery of morphology can be clarified by focusing on the acquisition process rather than the end-point. This has been our aim in this paper. Our MAL is a travesty of natural language, but at least we know the types and tokens in the learners' language evidence, and there is no need to speculate or argue about extrapolations from corpus data or assumptions about registers.

Human learning of this MAL inflectional morphology quickly culminates in a state where, as with natural language, frequency and regularity have interactive effects on performance. But, as we chart acquisition, it is clear that this interaction need not imply complex dual-mechanisms of processing. Rather, it simply reflects the asymptotes expected from the power law of practice, a simple associative law of learning. Thus, we have shown that one of the most frequently introduced arguments for the necessity of a dual-mechanism approach, a frequency effect for irregulars and the absence of such an effect for regulars, is not a good argument at all. Furthermore, we have demonstrated that a simple connectionist model, as an implementation of associative learning provided with the same language evidence, accurately simulates the human acquisition data.

But how is the power law instantiated in human and connectionist systems, and what is being associated in the acquisition of inflectional

morphosyntax? The power law of learning in human performance has been interpreted as resulting from basic associative mechanisms involving the formation of new chunks, and the effects of frequency on the accessibility of these representations (Newell, 1990, Newell & Rosenbloom, 1981). Anderson and Schooler (1991) suggest that memory (both as its behavioural expression in error rate and latency and as its neural expression in LTP) displays properties such as the power law of learning because these properties reflect an optimal response to the environment where the probability of an item occurring at any particular time is a power function of its past frequency of occurrence. Neural activation, which controls behaviour, reflects the probability of an item occurring in the environment; thus, the neural processes are designed to adapt behaviour to the statistical properties of the environment (Anderson, 1993). Connectionist systems are designed to do the same thing (Chater, 1995).

In our simplified account of inflectional morphology where phonological factors are put to one side, the relevant units for chunking are the stem forms and the plural affixes. From an associative perspective, regularity and frequency are essentially the same factor under different names. The first meaning of “regular” in the *Pocket Oxford Dictionary* involves “habitual, constant” acts, a definition in terms of statistical frequencies, consistency, and descriptive generalisation; the second stresses “conforming to a rule or principle”. We need to disentangle these senses (see Sharwood-Smith, 1994 and Lima, Corrigan, & Iverson, 1994 for conceptual analysis of “rules of language”). Whether regular morphology is generated according to a rule or not, it is certainly the case for English and the MAL under study here (and generally it is the default, if not the universal case—we will return to this matter later) that regular affixes are more habitual or frequent. And, as demonstrated in Fig. 5, the power law of practice entails that an effect of a constant increment of regularity (in its frequency sense) is much more apparent at low than at high frequencies of practice.

Although it is a general principle, the degree to which it applies depends on a range of factors, including: (a) the exponent of the power function, (b) the particular level of experience attained, and thus the placement of comparison points on the learning curve, and (c) the degree to which frequency and regularity are additive or multiplicative. In the present experiment, a fivefold increase in the frequency of the regular items results in a ($5 \times$ the number of regular items) increase in use of the regular affix; a fivefold increase in the frequency of an irregular item results in merely a fivefold increase in the use of the irregular affix. Thus, frequency and regularity are interactive rather than additive. But even if we allow for interaction, the function still results in greater regularity effects for low frequency items—just as, for example, the power function

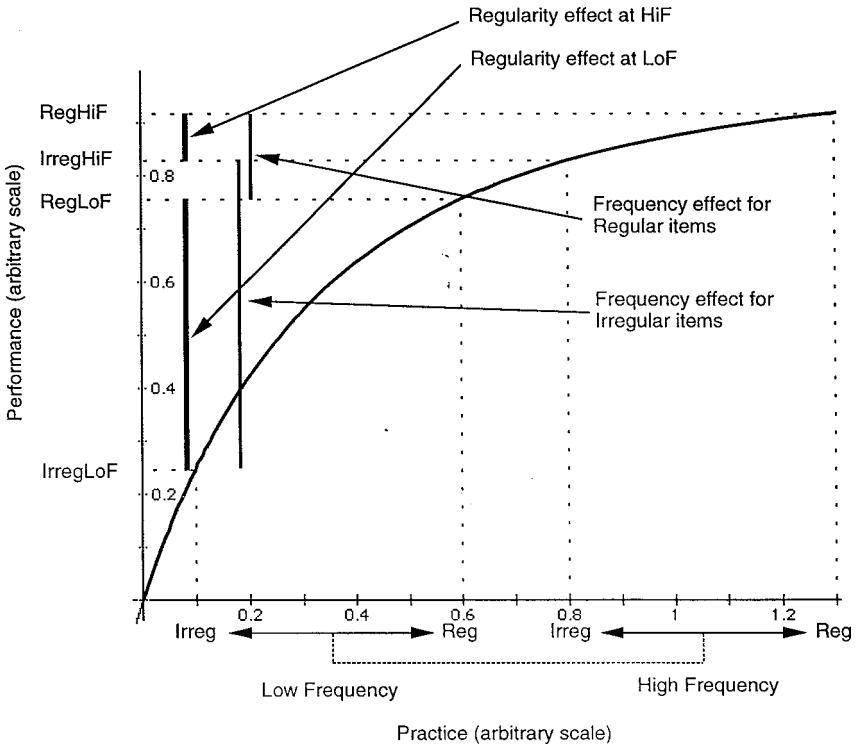


FIG. 5. A frequency by regularity interaction arising from additive contributions of regularity (solid horizontal arrows) and frequency (dotted horizontal arrows) inputting into an asymptoting power function. Notice in particular the solid vertical bars measuring out the large regularity effect at low frequencies, and the much smaller one at high frequencies. (Adapted from Plaut, McClelland, Seidenberg, & Patterson, 1994.)

$$y = 1 - x^{-2}$$

asymptotes, so does any power function

$$y = 1 - (x^n)^{-2}$$

where $n > 0$; the shape remains the same, albeit stretched or condensed along the horizontal axis. Thus all associative accounts of morphology, whether they stress the importance of type or token frequency (Bybee, 1995) in the determination of statistical regularity, imply a frequency by regularity interaction in performance.

Plaut et al. (1996) analyse the operation of connectionist networks in the particular quasi-regular domain of spelling-sound consistency in reading to demonstrate how the frequency by regularity interaction is a direct

consequence of the nonlinearity, adaptivity, and distributed representation properties of learning and representation in PDP networks. In what follows, we will minimally rephrase their analysis as it applies to the quasi-regular domain of inflectional morphology. In a connectionist network, the weight changes induced by an input/output pattern (I/OP) on any training epoch serve to reduce the error on that I/OP. The frequency of the I/OP (and the units it involves) is reflected in how often it is presented to the network. Thus word frequency directly amplifies weight changes that are helpful to the I/OP itself. Consistency of the morphological inflections of two stems is reflected in the similarity of affix units that are co-activated in their I/OPs. Furthermore, two inputs will induce similar weight changes to the extent that they activate similar units. In our MAL, as an extreme case, consistent forms all activate the same affix unit, irregular ones each activate a different idiosyncratic affix. Given that the weight changes that are induced by each I/OP are superimposed on the weight changes for all other I/OPs, an I/OP will tend to be helped by the weight changes for I/OPs whose input/output mappings are consistent with its own, and hindered by the weight changes for inconsistent I/OPs. Thus, frequency and consistency sum because they both arise from similar weight changes that are simply added together during training. The weight changes result in corresponding increases in the summed input to output units that should be active, and decreases to the summed units that should be inactive. However, due to the non-linearity of the input-output function of units, these changes do not produce directly proportionate reduction of error. Rather, as the magnitude of the summed input to output units increases, their states gradually asymptote towards 1.0—a given increase in the summed input to a unit yields progressively smaller decrements in error over the course of training. Thus, although frequency and regularity-as-consistency each contribute to the weights, and hence to the summed input to units, their effect on error is subjected to a gradual ceiling effect as unit states are driven towards extremal values.

Thus, a connectionist associative account of simple morphosyntax as it is embodied in our MAL holds that learning involves associating input patterns representing single or plural concepts with stem and affix lemmas across a large distributed network. Frequency of experience increases the strength of the appropriate I/O associations. Regularity effects stem from consistency, the consistent items all involve pairings between plurality and the regular lemma, and thus regularity is frequency by another name. The network sums and abstracts these consistencies, but it does so using non-linear unit input-output functions, thereby resulting in the frequency by regularity interaction. Networks are not simple competitive chunking or Markov chaining mechanisms working on surface form. Their massively distributed nature allows the emergence of more abstract internal representations. We have argued that this analysis accounts for the human

acquisition data of simple MAL morphosyntax quite well. We believe that the acquisition of natural language morphosyntax, where there are many additional factors of different phonological consistencies (of the type, for example, where the neighbours *sink*, *drink*, and *stink* are irregular in their past tenses but all behave in the same *-ank* way), are equally conducive to the principles of this type of account, although, as illustrated in grander simulation enterprises (Cottrell & Plunkett, 1994; Daugherty & Seidenberg, 1994; MacWhinney & Leinbach, 1991; Marchman, 1993; Plunkett & Marchman, 1993) the complexity of interaction of the factors that are there in the language evidence leads to much more complex developmental outcomes. Our role here has been to study human acquisition under precisely known circumstances and to demonstrate just how well a connectionist associative account can simulate these data.

A simple regularity=consistency account of this type will have difficulty if the "regular" or "default" case is not the most frequent case in a natural language. Although there is agreement for English past tense, and for morphology more generally, that the default case is more frequent, there may be exceptions. Marcus et al. (1995) argue that while the German particle *-t* applies to a much smaller percentage of verbs than its English counterpart, and the German plural *-s* applies only to a small percentage of nouns, nevertheless these affixes behave as defaults in the language. These default suffixations in German could thus pose a problem for statistical or connectionist accounts of the acquisition of the more frequent patterns as default since they may not be due to a large number of regular words reinforcing a pattern in associative memory (Prasada & Pinker, 1993). However, this is still a matter of some debate. Bybee (1995) suggests that a more reasonable method of counting German particle type frequency does show the default (or "productive") process to have the highest type frequency. She also argues that to a large extent the productivity patterns of German plurals also reflect their type frequency. Nakisa and Hahn (1996) and Plunkett and Nakisa (in press) show that generalisation to unseen or novel forms in German and Arabic (where there have also been claims for a minority default) is more accurately predicted by their phonological similarity to existing forms in the language (properly represented for type and token frequency), rather than by the operation of a default rule. Finally, Hare, Elman, and Daugherty (1995) demonstrate that multilayered networks can develop a default category even in the absence of superior type frequency, as long as the non-default classes are well defined and narrowly defined so that they serve as strong prototypes for analogising to novel forms. In such cases the area outside these well-defined attractor basins can constitute a potential default (see also, Plunkett & Marchman, 1991).

In the original hybrid model, irregulars were stored and accessed from rote memory. Pinker and Prince (1994, p. 326) modified this part of the

model, arguing that, since rote memory could not account (a) for similarities between the morphological base and irregular forms (e.g. *swing–swung*), (b) for similarity within sets of base forms undergoing similar processes (e.g. *sing–sang*, *ring–rang*, *spring–sprang*), or (c) for the kind of semi-productivity shown when children produce errors such as *bring–brang* or *swing–swang*, the memory system underlying such productions must be associative and dynamic, somewhat as connectionism portrays it. Yet to account for data such as the frequency/regularity interaction, this revised hybrid model still holds that regular forms are rule-governed. But a purely rule-based account of regulars cannot explain false friends effects where regular inconsistent items (e.g. *bake–baked* is similar in rhyme to neighbours *make–made* and *take–took*, which have inconsistent past tenses) are produced more slowly than entirely regular ones (Daugherty & Seidenberg, 1994; Seidenberg & Bruck, 1990), or frequency effects on regular forms (Oetting & Rice, 1993; Stemberger & MacWhinney, 1986). Unlike connectionist models, a rule-based account of regulars cannot explain these aspects of the human data. Nor is the regularity/frequency interaction any reason to reject connectionist accounts of morphosyntax in favour of a hybrid model.

REFERENCES

- Anderson, J.R. (1982). Acquisition of cognitive skill. *Psychological Review*, 89, 369–406.
- Anderson, J.R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Anderson, J.R., & Schooler, L.J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.
- Beck, M. (1995). Tracking down the source of NS–NNS differences in syntactic competence. Unpublished manuscript, University of North Texas.
- Bellugi, U., Bihle, A., Jernigan, D., Trauner, D., & Dougherty, S. (1990). Neuropsychological, neurological, and neuroanatomical profile of Williams Syndrome. *American Journal of Medical Genetics*, 6, 115–125.
- Braine, M.D.S., Brody, R.E., Brooks, P.J., Sudhalter, V., Ross, J.A., Catalano, L., & Fisch, S.M. (1990). Exploring language acquisition in children with a miniature artificial language: Effects of item and pattern frequency, arbitrary subclasses, and correction. *Journal of Memory and Language*, 29, 591–610.
- Broeder, P., & Plunkett, K. (1994). Connectionism and second language acquisition. In N. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 421–454). London: Academic Press.
- Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes*, 10, 425–455.
- Chater, N. (1995). Neural networks: The new statistical models of mind. In J.P. Levy, D. Bairaktaris, J.A. Bullinaria, & P. Cairns (Eds), *Connectionist models of memory and language*. London: UCL Press.
- Cohen, J.D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments and Computers*, 25, 257–271.
- Cottrell, G., & Plunkett, K. (1994). Acquiring the mapping from meaning to sounds. *Connection Science*, 6, 379–412.

- Daugherty, K.G., & Seidenberg, M.S. (1992). Rules or connections? The past tense revisited. In *Proceedings of the 14th annual conference of the Cognitive Science Society* (pp. 259–264). Pittsburgh, PA: Cognitive Science Society.
- Daugherty, K.G., & Seidenberg, M.S. (1994). Beyond rules and exceptions: A connectionist approach to inflectional morphology. In S.D. Lima, R.L. Corrigan, & G.K. Iverson (Eds), *The reality of linguistic rules* (pp. 353–388). Amsterdam: John Benjamins.
- DeKeyser, R. (1997). Beyond explicit rule learning: Automatizing second language morphosyntax. *Studies in Second Language Acquisition*, 19, 195–222.
- Ellis, N.C. (1996). Sequencing in SLA: Phonological memory, chunking, and points of order. *Studies in Second Language Acquisition*, 18, 91–126.
- Eubank, L., & Gregg, K.R. (1995). “Et in Amygdala Ego”? UG, (S)LA, and neurobiology. *Studies in Second Language Acquisition*, 17, 35–58.
- Hare, M., Elman, J.L., & Daugherty, K.G. (1995). Default generalisation in connectionist networks. *Language and Cognitive Processes*, 10, 601–630.
- Jung, J. (1971). *The experimenter's dilemma*. New York: Harper & Row.
- Kirsner, K. (1994). Implicit processes in second language learning. In N. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 283–312). London: Academic Press.
- Lachter, J., & Bever, T. (1988). The relation between linguistic structure and associative theories of language learning: A constructive critique of some connectionist learning models. *Cognition*, 28, 195–247.
- Lima, S.D., Corrigan, R.L., & Iverson, G.K. (Eds). (1994). *The reality of linguistic rules*. Amsterdam: John Benjamins.
- MacWhinney, B. (1983). Miniature language systems as tests of use of universal operating principles in second-language learning by children and adults. *Journal of Psycholinguistic Research*, 12, 467–478.
- MacWhinney, B. (1994). The dinosaurs and the ring. In S.D. Lima, R.L. Corrigan, & G.K. Iverson (Eds). *The reality of linguistic rules* (pp. 283–320). Amsterdam: John Benjamins.
- MacWhinney, B., & Leinbach, J. (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition*, 40, 121–157.
- Marchman, V.A. (1993). Constraints on plasticity in a connectionist model of the English past tense. *Journal of Cognitive Neuroscience*, 5, 215–234.
- Marcus, G.F., Brinkmann, U., Clahsen, H., Wiese, R., & Pinker, S. (1995). German inflection: The exception that proves the rule. *Cognitive Psychology*, 29, 198–256.
- McLaughlin, B. (1980). On the use of miniature artificial languages in second-language research. *Applied Psycholinguistics*, 1, 357–369.
- Moeser, S.D., & Bregman, A.S. (1972). The role of reference in the acquisition of a miniature artificial language. *Journal of Verbal Learning and Verbal Behavior*, 11, 759–769.
- Morgan, J.L., Meier, R.P., & Newport, E.L. (1987). Structural packaging in the input to language learning: Contributions of prosodic and morphological marking of phrases to the acquisition of language. *Cognitive Psychology*, 19, 498–550.
- Morgan, J.L., & Newport, E.L. (1981). The role of constituent structure in the induction of an artificial language. *Journal of Verbal Learning and Verbal Behavior*, 20, 67–85.
- Morton, J. (1979). Facilitation in word recognition: Experiments causing change in the logogen model. In P.A. Kollers, M.E. Wroldstad, & M. Bouma (Eds), *Processing of visible language* (pp. 259–268). New York: Plenum.
- Nakisa, R., & Hahn, U. (1996). Where defaults don't help: The case of the German plural system. In *Proceedings of the 18th annual conference of the Cognitive Science Society* (pp. 177–182). Hillsdale, NJ: Lawrence Erlbaum Associates Inc.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Rosenbloom, P. (1981). Mechanisms of skill acquisition and the law of practice. In J.R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates Inc.

- Oetting, J.B., & Rice, M.L. (1993). Plural acquisition in children with specific language impairment. *Journal of Speech and Hearing Research*, 36, 1236–1248.
- Paivio, A. (1986). *Mental representations. A dual coding approach*. Oxford, UK: Oxford University Press.
- Palermo, D.S., & Howe, H.E. (1970). An experimental analogy to the learning of past-tense inflection rules. *Journal of Verbal Learning and Verbal Behavior*, 9, 410–416.
- Pinker, S. (1991). Rules of language. *Science*, 253, 530–535.
- Pinker, S., & Prince, A. (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition*, 29, 195–247.
- Pinker, S., & Prince, A. (1994). Regular and irregular morphology and the psychological status of rules of grammar. In S.D. Lima, R.L. Corrigan, & G.K. Iverson (Eds), *The reality of linguistic rules* (pp. 321–351). Amsterdam: John Benjamins.
- Plaut, D.C., McClelland, J.L., Seidenberg, M.S., & Patterson, K.E. (1996). Understanding normal and impaired word reading: Computational principles in quasi-regular domains. *Psychological Review*, 103, 56–115.
- Plunkett, K., & Marchman, V. (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition*, 38, 3–102.
- Plunkett, K., & Marchman, V. (1993). From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition*, 48, 21–69.
- Plunkett, K., & Nakisa, R.C. (in press). A connectionist model of Arabic plural system. *Language and Cognitive Processes*.
- Prasada, S., & Pinker, S. (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes*, 8, 1–56.
- Prasada, S., Pinker, S., & Snyder, W. (1990). Some evidence that irregular forms are retrieved from memory but regular forms are rule-governed. Paper presented at the 31st meeting of the Psychonomic Society. New Orleans, November.
- Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representations by back propagation. In D.E. Rumelhart & J.L. McClelland (Ed), *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
- Rumelhart, D., & McClelland, J. (1986). On learning the past tense of English verbs. In D.E. Rumelhart & J.L. McClelland (Eds), *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models* (pp. 272–326). Cambridge, MA: MIT Press.
- Seidenberg, M.S., & Bruck, M. (1990). Consistency effects in the generation of past tense morphology. Paper presented at the 31st meeting of the Psychonomic Society. New Orleans, November.
- Seidenberg, M.S., Waters, G.S., Barnes, M.A., & Tanenhaus, M.K. (1984). When does irregular spelling or pronunciation influence word recognition? *Journal of Verbal Learning and Verbal Behavior*, 23, 383–404.
- Sharwood Smith, M.A. (1994). The unruly world of language. In N. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 33–44). London: Academic Press.
- Snodgrass, J.G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, image agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 174–215.
- Stemberger, J.P., & MacWhinney, B. (1986). Frequency and the lexical storage of regularly inflected forms. *Memory and Cognition*, 14, 17–26.
- Winter, B., & Reber, A.S. (1994). Implicit learning and the acquisition of natural languages. In N. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 115–146). London: Academic Press.
- Yang L.R., & Givón, T. (1997). Benefits and drawbacks of controlled laboratory studies of second language acquisition: The Keck second language learning project. *Studies in Second Language Acquisition*, 19, 173–194.

APPENDIX

The Word-forms of the Artificial Language

<i>Picture</i>	<i>Stem</i>	<i>Plural Form</i>	<i>Frequency</i>	<i>Regularity</i>
car	garth	bugarth	5	R
bed	pid	bupid	1	R
lamp	lant	bulant	5	R
table	tib	butid	1	R
plane	poon	bupoon	5	R
ball	prill	buprill	1	R
train	dram	budram	5	R
house	hize	buhize	1	R
book	bisk	bubisk	5	R
broom	breen	bubreen	1	R
phone	feem	gofeem	5	I
umbrella	brol	gubrol	1	I
chair	charp	zecharp	5	I
horse	naig	zonaig	1	I
monkey	chonk	nuchonk	5	I
dog	woop	niwoop	1	I
elephant	fant	vefant	5	I
scissors	zoze	vuzoze	1	I
kite	kag	rekag	5	I
fish	pisc	ropisc	1	I